Chapter 6
Conditionals

John Cantwell

Abstract Conditional constructions – constructs of the form If A, then B – have for over a century been subject to intense study in a wide variety of philosophical areas, as well as outside of philosophy. One important reason is that such constructs allow one to encode *connections* and *dependencies*, be they causal, epistemic, conceptual, or metaphysical. This chapter briefly outlines some of the main formal models that have been employed to analyze such constructs, as well as their philosophical motivation.

6.1 Background

The conditional construction – here I include such constructions as *If A then B*, *B even if A*, *B only if A*, and so on – is a small unassuming construction that for decades (in some cases centuries) has attracted massive interest from philosophers, logicians, linguists, computer scientists and cognitive psychologists. The huge interest in this small construction can be traced to two circumstances. First, the conditional is one of our primary vehicles for talking about, describing and representing *connections* and *dependencies*, be the connections and dependencies causal, conceptual, metaphysical, epistemic, or logico-semantic. Second, the problems one has in accounting for the meaning of the conditional to a large extent overlaps with the problems one has in accounting for the nature of these connections and dependencies. So the study of the conditional is one pathway into a large body of issues with ramifications far beyond the seemingly minor issue of the semantics of a small unassuming construction. Importantly, the formal methods that have been introduced in order to deal with the problems posed by conditionals have proved to be useful and illuminating in a range of other areas. Not surprisingly, the area – with its connections to many of the deepest problems of philosophy – is rife with controversy.

J. Cantwell (✉)
Division of Philosophy, Royal Institute of Technology (KTH), Stockholm, Sweden
e-mail: cantwell@kth.se

S. O. Hansson, V. F. Hendricks (eds.), *Introduction to Formal Philosophy*, Springer Undergraduate Texts in Philosophy, https://doi.org/10.1007/978-3-319-77434-3_6

Consider the following conditionals:

(1) If x is even and greater than 2, then x is not prime.
(2) If the conditions of the Versaille treaty had not been so severe, there would have been no WWII.
(3) If Shakespeare didn't write Hamlet someone else did.

The first of these can be used to express that *even* and *prime* are conceptually related properties; the second conditional can be used to express that there is a causal connection between the conditions set up in the Versaille treaty and the breakout of WWII; the third conditional can be be used to express that there is some degree of epistemic independence between one's conviction that *Shakespeare* wrote Hamlet and one's conviction that *someone* wrote Hamlet (e.g. some of the evidence that one has for the latter, is not also evidence that one has for the former). Three conditionals that can be used to express three very different kinds of connections and dependencies.

A substantial portion of the literature on conditionals deals with the semantics of conditionals in natural language. Various fundamental semantic questions have been addressed. Is there at some deeper semantic level only one type of conditional or are there different semantic kinds of conditionals? Compare the difference between (3) above, which is in the grammatically *indicative* mood, and the grammatically *subjunctive*:

(4) If Shakespeare hadn't written Hamlet someone else would have.

Clearly they don't have the same meaning, and this has convinced many that that there are different semantic kinds of conditionals. Further questions: Does the semantic value of a conditional depend on the (conceptual/causal/epistemic) connections and dependencies it is used to express, or does the conditional express the connections and dependencies that it does through the pragmatics of assertion? Can conditionals have truth values and, if so, do they always have truth values (or do they express 'gappy' propositions, propositions that can lack truth value)? Is the truth value of a conditional (if it has one) a function of the truth values of its constituent sentences or is the conditional semantically intensional? Note that some of these questions may have different answers for different kinds of conditionals.

Many, however, have approached the study of conditionals from another direction. Seeing that conditionals are linguistic vehicles for expressing connections and dependencies one may ask what kind of connections and dependencies – perhaps encoded in some underlying structure – they can be used to express. In such studies felicity to natural language usage is of secondary importance (although the link and appeal to linguistic intuitions is seldom abandoned altogether), instead more general structural phenomena are investigated.

One important phenomenon is *defeasibility* (sometimes referred to as *failure of antecedent strengthening* or *nonmonotonicity*). From the fact that one accepts $A \rightarrow C$ (I here use $\rightarrow$ as a generic conditional) it does not follow that one thereby should accept $(A \wedge B) \rightarrow C$. For instance, from the fact that one accepts:

(5) If the match is struck it will light.

it does not follow that one should accept:

(6) If the match is *submerged in water* and struck it will light.

The phenomenon of defeasibility reflects the fact that connections and dependencies often do not hold unconditionally or by necessity; often they depend on things being as they *normally* are, or on contingencies that just happen– as far as we know – to be the case or on other things *being equal* (*ceteris paribus*). A fundamental problem is that typically it is *impossible* to spell out in full detail how things normally are, what contingencies are necessary and sufficient or what the 'other things' are that are to be considered 'equal'. The expressive power of conditionals to a large extent derives from the fact that they implicitly invoke dependencies that are impossible to spell out. As nearly all discourse outside the realm of mathematics (and other aprioristic disciplines) deals with such defeasible connections and dependencies, it is of the first importance to understand how conditionals do the magic of depending on what cannot be spelled out. Note that we will make no progress here by merely *mentioning* that a conditional depends on that things are as they normally or on other things being equal. For from the fact that one accepts:

(7) If the match is struck it will, *other things being equal*, light.

it still does not follow that one should accept:

(8) If the match is submerged in water and struck it will, *other things being equal*, light.

The phenomenon of defeasibility is a core feature of the conditional itself, and the development of formal methods for analysing this phenomenon has been of central interest in the literature on conditionals.

In this paper I shall briefly discuss some of the main approaches to the analysis of conditionals, outline the formal structures that have made the analyses possible, and briefly indicate their philosophical underpinnings, with a particular eye towards how they account for how conditionals can be used to express connections and dependencies. I will consider two basic kinds of analyses: those that assign truth conditions to the conditional, and those that instead assign acceptance conditions (using the *Ramsey Test*).[1]

[1]The field is far too wide to enable an exhaustive survey in these few pages and so the present overview largely reflects the author's own interests and prejudices. Many important issues are ignored or treated only in passing and countless important contributions will not be credited. References given reflect (but by no means exhaust) works of seminal importance, works that give a more thorough overview of the issues, as well as work that may point towards interesting new developments.

6.2 Conditionals with Truth Values

6.2.1 Formal Preliminaries

Assume a language with atoms $p, q, r, \ldots$ closed under $\neg$ (negation), $\wedge$ (conjunction) and $\vee$ (disjunction). The language will subsequently be extended with the generic conditional $\rightarrow$. Non-atomic sentences will be denoted A, B, C, ...

The semantics for this language will be given relative a set U of *states*; these can be thought of as assignments of truth values to the atoms or as *possible worlds*, or as pairs of possible worlds and moments in time, what is crucial is that each state determines the truth values of the atomic sentences.

1. $u \models p$ if and only if p is true at u.
2. $u \models \neg A$ if and only if it is not the case that $u \models A$.
3. $u \models A \wedge B$ if and only if $u \models A$ and $u \models B$.
4. $u \models A \vee B$ if and only if $u \models A$ or $u \models B$.

In the form of truth-tables, at any given state u:

A	$\neg A$
t	f
f	t

A	B	$A \wedge B$
t	t	t
t	f	f
f	t	f
f	f	*f*

A	B	$A \vee B$
t	t	t
t	f	t
f	t	t
f	f	f

A sentence B is a *consequence* of $A_1, \ldots, A_n$, in symbols $A_1, \ldots, A_n \models B$ if and only if $u \models B$ whenever $u \models A_1, \ldots, u \models A_n$ (that is, if and only if B is true whenever each A_i is true).

6.2.2 The Material Conditional

The *material conditional*, here written $A \supset B$, is typically the first fully analysed conditional that students of logic encounter. In its classical form it is a truth-functional connective with interpreted according to the truth-table:

A	B	$A \supset B$
t	t	t
t	f	f
f	t	t
f	f	t

The interpretation has the virtue of simplicity (truth-tables belong to the simplest of formal structures), furthermore it can be derived from seemingly compelling logical principles; for instance, one can show that it is the interpretation that $\supset$ must have if it is to be truth-functional and satisfy (given the classical interpretation of negation):

$$A, A \supset B \models B. \text{ (Modus ponens)}$$

$$\neg B, A \supset B \models \neg A. \text{ (Modus tollens)}$$

$$\text{If } A \models B, \text{ then } \models A \supset B.$$

It is clear from the semantics of the material conditional that its truth value does not depend in any interesting way on connections and dependencies between antecedents and consequents: the material conditional is only sensitive to their truth value. As a result, if we take, say, the indicative conditional of natural language to have the semantics of the material conditional, semantics alone does very little to explain how and when such conditionals are asserted and denied. On the basis of semantics alone one would predict that both of the following conditionals would be generally accepted:

(9) If Shakespeare didn't write Hamlet, then his grandmother did.
(10) Even if someone other than Shakespeare wrote Hamlet, Shakespeare wrote Hamlet.

For if Shakespeare wrote Hamlet, then both these conditionals are true (as their antecedent is false). Yet those who believe that Shakespeare wrote Hamlet are not in general inclined to accept the conditionals. Indeed most would be inclined to *reject* these conditionals; so there is a deep and disturbing mismatch between the truth values of the conditionals and speakers' inclinations to use them. Such systematic mismatch is sometimes labeled 'paradoxes' of the material conditional and can be traced back to various logical properties such as:

$$\neg A \models A \supset B.$$

$$B \models A \supset B.$$

If the semantics of the material conditional is to have any credibility as a semantics of the conditional in natural language (specifically: the indicative conditional) then most of the explanatory work – including how such conditionals are used to express connections and dependencies – must be relegated elsewhere: to pragmatics. For instance, one can hold that the reason why one is not inclined to accept (9) is the same as the reason why one is not inclined to assert

(11) Either Shakespeare wrote Hamlet or his grandmother did.

even though one believes that it is true (due to one of the disjuncts being true): it would be misleading and would convey less information than the simple "Shakespeare wrote Hamlet". This was the strategy proposed by Grice [12] and has subsequently been defended by Lewis [22] and elaborated by Jackson [13]. However, many have come to the conclusion that the discrepancy between the semantics of the material conditional and the way in which indicative conditionals are used is too great: the semantics lacks credibility.

Some have argued (e.g. [5, 7]) that the most blatant collisions between semantics and pragmatics (e.g. cases where one is inclined to reject a conditional that one, allegedly, believes is true) can be avoided by allowing conditionals to have ‘gappy’ truth conditions:

A	B	$A \to B$
t	t	t
t	f	f
f	t	–
f	f	–

That is, the conditional is taken to lack truth value when the antecedent is false. This still leaves much of the explanatory work to pragmatics (e.g. the way the conditional is used to express connections and dependencies) but at least does not force pragmatics to explain why true conditionals are rejected. The semantics also has support from the psychological literature on how people assess conditionals. Nevertheless, the account is still viewed with scepticism, mainly regarding the intelligibility of truth-value gaps and how such gaps are to be accommodated in a wider story where conditionals embed in more complex sentences (see [23, 24]).

6.2.3 The Strict Conditional

An example of a conditional that semantically reflects a stronger connection between the antecedent and the consequent is the *strict* conditional, with the truth-conditions:

$u \models A \to B$ if and only if $v \models B$ for every state v such that $v \models A$.

The strict conditional is not supposed to reflect any particular construction in natural language, but it is a nice simple example of how a language can contain conditionals that reflect interesting connections between their antecedents and consequents, in this case: the connection of semantic consequence (for with the present truth-conditions, $A \to B$ is true if and only if B is a semantic consequence of A). If the English indicative conditional had these truth conditions then

(12) If you are a bachelor then you are not married,

would be true, while

(13) If Jim was run over by a truck, he died,

would be false (even if Jim was run over by a truck and died).

Notably, as semantic consequence involves the preservation of a property (the property of being true at a state), the strict conditional is not defeasible, that is, we have the property:

$$A \to C \models (A \land B) \to C.$$

6.2.4 Ontic Selection Functions: Counterfactuals

An explosion of interest in the formal structures used to represent conditionals came with the work of Stalnaker [29, 31] and Lewis [19]. In different ways they introduced the idea of a selection function γ that for each state u and each sentence A picks out a sub-set of the states in which A is true. Given such a selection function one can give truth conditions for a conditional as follows:

$$u \models A \to B \text{ if and only if } v \models B \text{ for every } v \text{ in } \gamma(u, A).$$

A selection function gives rise to many degrees of freedom both in terms of abstract structural properties and in terms of substantive interpretations. Common examples of structural constraints are[2]:

If v is in $\gamma(u, A)$, then $v \models A$.
If $u \models A$, then $\gamma(u, A)$ is the unit set $\{u\}$. (Centering)
If $v \models B$ for each v in $\gamma(u, A)$, then $\gamma(u, A) = \gamma(u, A \land B)$.

In these cases each structural constraint gives rise to a logical property:

$$\models A \to A.$$
$$A, A \to B \models B.$$
$$A \to B, A \to C \models (A \land B) \to C.$$

Notably, the semantics allows for *defeasibility*, that is, we do not in general have:

$$A \to C \models (A \land B) \to C.$$

For instance, consider the graphical representation in Fig. 6.1 of the states in which A, B and C are true and the areas selected by the selection function:
From this picture it is clear that we can have $u \models A \to C$ although we do not have $u \models (A \land B) \to C$.

Relationships between structural properties of the selection function and logical properties of the conditional are, of course, interesting in their own right. But it

[2] Arló-Costa [3] presents a thorough overview of the logic of conditionals and how they relate to structural conditions.

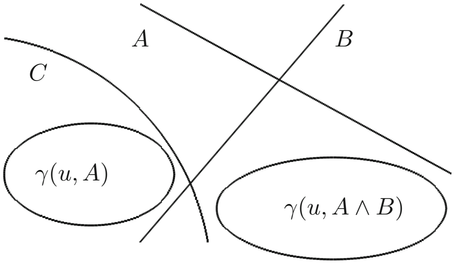

Fig. 6.1 $\gamma(u, A)$ is a subset of the C-states while $\gamma(u, A \wedge B)$ is not

has been the prospect that selection functions can be used to represent fundamental structures – the kind of structures that underlie the connections and dependencies that we express by natural language conditionals – that has been the main philosophical driving force for investigating selection-function semantics for conditionals.

By far the most influential interpretation is due to David Lewis (e.g. [19, 21]) who suggested that the selection function $\gamma(u, A)$ selects those A-worlds (on his account states are possible worlds) that are *most similar* to u; that is, the selection function is based on a similarity relation $v \leq_u w$ between possible worlds ($v \leq_u w$ holds if v is no less similar to u than is w). The similarity relation is standardly taken to be *reflexive*, *transitive* and *complete*[3] and to satisfy some proviso – *the limit assumption* – to ensure that in a given set of worlds there is at least one world that is *most similar* to the target world, that is one needs to guarantee that there are no infinitely descending chains of similarity.[4]

Similarity relations between worlds can be thought of as providing a substantive interpretation of the *ceteris paribus*-intuition in the evaluation of conditionals: a conditional is true if it is the case that if the antecedent were true, and *as much as possible* (this is where the similarity relation kicks in) remained the same, then the consequent would be true.

The introduction of a similarity relation shifts the focus to the question: *Similar in what way?* There are various ways of approaching this question. Lewis, who followed a tradition broadly deriving from David Hume, took special interest in similarity criteria that would allow for a reductive analysis of causal relations in terms of the similarity relation (and so for a reductive analysis of causal relations in terms of counterfactuals).

Counterfactual analyses of causality have been criticized (e.g. [32]), but this is not the place to evaluate Lewis' theory of causality. What is clear however is that a Lewis-style similarity semantics provides a robust and convincing model (even if one may be sceptical about the underlying metaphysical assumptions) for the

[3]**Reflexivity:** $v \leq_u v$. **Transitivity:** If $v \leq_u w$ and $w \leq_u z$, then $v \leq_u z$. **Completeness:** Either $v \leq_u w$ or $w \leq_u v$.

[4]The limit assumption is not, strictly speaking, necessary, but if one omits this constraint the semantic clause becomes more complex.

analysis of *counterfactual* conditionals, a class of conditionals that largely coincides with conditionals in the *subjunctive* mood:

(14) If the match had been struck, it would have lit.
(15) The Vietnam war would have escalated even if Kennedy had not been murdered.

As we typically take the truth values of such conditionals to depend on underlying causal connections and dependencies, this provides strong support for the idea that similarity relations can be used to encode important aspects of the causal structure of the world. Accordingly, the Lewis-Stalnaker analysis remains the dominant paradigm in the semantics of counterfactual conditionals, one where the *truth-value* of a conditional is sensitive to underlying causal connections and dependencies between antecedents and consequents.

Within the more linguistically oriented study of the semantics of conditionals, the dominant tradition is to take conditionals to have an underlying Lewis-Stalnaker-style semantic structure. Some of the most influential work here is due to Angelika Kratzer (see e.g. [14, 15], see collection in [16]). Kratzer combines a Lewis-style semantics with a 'restrictor' analysis of conditionals: the antecedent of a conditional is taken to restrict the space of possibilities relative to which the consequent is evaluated. This becomes particularly important when the consequent of a conditional contains a modality of some sort, as in:

(16) If the die is thrown lands on an even number, the probability that it will show a six is $1/3$.
(17) If you kill him, you should kill him gently.

In the first case the antecedent constrains the possible outcomes that is relevant for the probability modality (for note that absent the antecedent, the probability that the result of a throw of a fair die will show a six is $1/6$). In the second case the antecedent constrains the space of possible actions to be considered in deliberating what one *should* do (and note that it is *not* the case that you should kill him gently, but *if* you kill him, then a gentle killing seems to be the most humane course of action).

6.3 Epistemic Interpretations

Epistemic interpretations of conditionals[5] cannot be discussed without mention of Frank Ramsey's [28] famous footnote:

> If two people are arguing 'if p will q?' and both are in doubt as to p, they are adding p hypothetically to their stock of knowledge and arguing on that basis

[5]Some key references here are Stalnaker [30], Adams [1], Edgington [9], Levi [18] and Bennett [4].

> about q. We can say that they are fixing their degrees of belief in q given p. If p turns out false, these degrees of belief are rendered void. If either party believes $\neg p$ for certain, the question ceases to mean anything to him except as a question about what follows from certain laws and hypotheses.

In the footnote Ramsey identifies the main feature of what has become known as the *Ramsey Test*: conditionals are evaluated on the basis of what one would take to hold on the hypothetical assumption that the antecedent is true. As hypothetical reasoning is a key way of exploring epistemic connections such as evidential relations and other structures of justification, the Ramsey Test points the way to explaining how conditionals can be used to express such connections.

For instance, say that you have quite convincing testimony from the butler that suggests that either the maid or the gardener commited the murder. So on supposing, hypothetically, that it wasn't the maid, you conclude (hypothetically) that it was the gardener; that is, you accept

(18) If the maid didn't do it, the gardener did.

On the other hand, on supposing that neither the maid nor the gardener did it, you come to the conclusion that there must be something wrong with the butler's testimony, indeed that would suggest that *he* is the culprit; so you accept:

(19) If neither the maid nor the gardener did it, the butler did it.

The fact that you accept both conditionals reveal something important about how you evaluate the situation and what counts as evidence for what (note, in particular, that the examples show that the epistemic interpretation allows for defeasible conditionals).

Epistemic interpretations of conditionals typically do not take the connections and dependencies expressed by the use of a conditional to reside in its truth-conditions. The epistemic connections and dependencies expressed by a conditional are not taken to be objective features of the world, but are rather features of one's current epistemic state.

There is wide agreement[6] that the epistemic intepretation provides a good analysis of stand-alone (non-embedded) *indicative* conditionals (conditionals in the indicative mood) like (3), (18), and (19) above.[7] From the point of view of meaning theoretical orthodoxy the big problem with the epistemic interpretation is that it doesn't provide truth-conditions for the conditional. This creates a problem both in accounting for its logic, as the semantic consequence relation is based on sentences taking a truth value, and in accounting for its interaction with other connectives like negation ($\neg$) and disjunction ($\vee$), as these are truth-functional. Some have thus sought to combine epistemic interpretations with truth-functional

[6]See [4] and [10] for extensive discussions and references, but compare [26] and [25] for putative counterexamples (e.g. [8] argues that McGee fails to establish a counterexample).

[7]Some have argued that counterfactual conditionals can be given an epistemic interpretation, but this is a matter of considerable controversy.

accounts and take the epistemic interpretation to spell out, in a systematic way, the pragmatics of indicative conditionals, thus maintaining the traditional semantic-pragmatics distinction (see the discussion in Sect. 6.2.2). Others have taken this to show that the indicative conditional is inherently different from other truth-conditional constructions and has no semantics proper.

6.3.1 The Ramsey Test and Logic

In its most abstract form, the epistemic interpretation relies on the notion of an epistemic state $\mathcal{E}$ and on a function $*$ that takes a sentence A and returns the epistemic state $\mathcal{E} * A$ that corresponds to the epistemic state in which it is hypothetically assumed that A. As neither the acceptance conditions of conditionals nor their logic are derivable from their truth-conditions, one also needs a separate account of acceptance conditions and of *epistemic* consequence.

Let $\mathcal{E}|\!\models A$ stand for *in the epistemic state $\mathcal{E}$ one is committed to accepting A* (some models, such as the probabilistic model, allow also for a notion of *degree* of acceptability, but this notion will not be covered here). With this in place one can state the Ramsey Test as follows:

$$\mathcal{E}|\!\models A \rightarrow B \text{if and only if } \mathcal{E} * A|\!\models B.$$

The Ramsey Test by itself gives no indication of the acceptance conditions for other kinds of sentences or of their logic. So let $A_1, \ldots, A_n|\!\models B$ stand for *B is an epistemic consequence of $A_1, \ldots, A_n$*. A reasonable minimal requirement is that as long as the sentences are not conditionals (nor contain conditionals) we have:

Semantic Closure If $A_1, \ldots, A_n \models B$, then $A_1, \ldots, A_n|\!\models B$.

A further reasonable requirement is:

Epistemic Closure If $\mathcal{E}|\!\models A_1, \ldots, \mathcal{E}|\!\models A_n$ implies $\mathcal{E}|\!\models B$ for every epistemic state $\mathcal{E}$, then $A_1, \ldots, A_n|\!\models B$.

If we allow for the converse direction then the epistemic consequence relation can be fully analyzed by acceptance conditions:

Reverse Epistemic Closure If $A_1, \ldots, A_n|\!\models B$, then $\mathcal{E}|\!\models A_1, \ldots, \mathcal{E}|\!\models A_n$ implies $\mathcal{E}|\!\models B$.

The logical properties of the conditional will to a large extent depend on properties of the $*$-operator. Here are two candidate properties (again, see Arló-Costa [3] for thorough overview):

Success $\mathcal{E} * A|\!\models A$.
Vacuity If $\mathcal{E}|\!\models A$, then $\mathcal{E} * A = E$.

These give rise to the properties:

$$\|\!\!\models A \to A.$$

$$A, A \to B \|\!\!\models B.$$

Another possible requirement is:

Iteration $\mathcal{E} * A * B = E * (A \wedge B)$.

This gives rise to the logical export-import-properties:

$$A \to (B \to C) \|\!\!\models (A \wedge B) \to C.$$

$$(A \wedge B) \to C \|\!\!\models A \to (B \to C).$$

Importantly, it is typically not assumed that $*$ satisfies *monotonicity*:

Monotonicity If $\mathcal{E} \|\!\!\models B$, then $\mathcal{E} * A \|\!\!\models B$.

Accordingly, the epistemic conditional can be defeasible: it is not in general the case that $A \to C \|\!\!\models (A \wedge B) \to C$.

Sometimes weaker properties than monotonicity are assumed, such as:

Preservation If $\mathcal{E} \|\!\!\models B$ and $\mathcal{E} \not\|\!\!\models \neg A$, then $\mathcal{E} * A \|\!\!\models B$.

Together with Logical Closure this entails:

If $\mathcal{E} \|\!\!\models A \supset B$ and $\mathcal{E} \not\|\!\!\models \neg A$, then $\mathcal{E} \|\!\!\models A \to B$.

This can be combined with the requirement:

If $\mathcal{E} * A \|\!\!\models B$ and $\mathcal{E} \not\|\!\!\models \neg A$, then $\mathcal{E} \|\!\!\models A \supset B$.

Together with Preservation and Logical Closure this entails that as long as one doesn't reject the antecedent of a conditional the acceptance conditions of the epistemic conditional coincides with acceptance conditions of the material conditional. This is important as the counterintuitive properties of material implication (as a model for the indicative conditional) mainly emerge in cases when the antecedent is believed to be false; the epistemic conditional thus keeps the 'good' parts of the material analysis of the conditional but avoids its problematic parts. For instance, in the absence of the monotonicity requirement we do not have:

$$\neg A \|\!\!\models A \to B.$$

$$B \|\!\!\models A \to B.$$

6.3.2 Formal Models

Two main types of formal structures are often used to model such epistemic states: models that assign *probabilities* to sentences and a broad range of models – closely related or identical to the models applied in the area of *belief revision* – that rank sentences according to their *entrenchment* (*plausibility*, etc.). Often (but not always) the latter kind of models use qualitative relations rather than numerical measures.

Probabilistic models (see in particular Adams [1]) follow a Bayesian tradition in which epistemic states are taken to be probability measures.[8] The revision operator $*$ in such a model takes a probability measure $\mathcal{E}_\mathrm{P}$ and a sentence A and returns a new probability measure $\mathcal{E}_\mathrm{P} * A$ by *conditionalisation*[9]:

$$(\mathcal{E}_\mathrm{P} * A)(B) = \mathcal{E}_\mathrm{P}(B \mid A) = \mathcal{E}_\mathrm{P}(A \wedge B)/\mathcal{E}_\mathrm{P}(A).$$

In a simple model of probabilistic acceptance one accepts all and only those non-conditional sentences that exceed some threshhold α ($.5 \leq \alpha \leq 1$), so that, when A is not a conditional:

$$\mathcal{E}_\mathrm{P} \Vdash A \text{ if and only if } \mathcal{E}_\mathrm{P}(A) > \alpha.$$

For the epistemic consequence relation there are different options; Adams [2] makes a case for the p-consequence relation (the uncertainty of the conclusion cannot be greater than the sum of the uncertainty of the premises):

$A_1, \ldots, A_n \Vdash B$ if and only if $(1-\mathcal{E}_\mathrm{P}(B)) \leq (1-\mathcal{E}_\mathrm{P}(A_1))+\cdots+(1-\mathcal{E}_\mathrm{P}(A_n))$, for all $\mathcal{E}_\mathrm{P}$.

These jointly ensure that Semantic Closure and Epistemic Closure are satisfied. The model ensures that we have a defeasible conditional that satisfies important logical properties like modus ponens and export-import.

As an example of a qualitative representation of a epistemic state one can consider an epistemic state to be represented as an epistemic selection function $\mathcal{E}_\gamma$ that for any sentence B that has truth conditions picks out the set $\mathcal{E}_\gamma(B)$ of *most plausible* B-states (the B-states, recall, are the states in U where B is true). So here again we have a selection function, but notice that while selection functions in the ontic models were relativized to the states (worlds) themselves, here they are relativized to *epistemic states*, making the evaluation procedure speaker-dependent. The epistemic state that results from making the hypothetical assumption that A,

[8] A probability measure is, as is standard, here taken to be a real-valued function P that take sentences as their arguments and satisfies (a) $0 \leq \mathrm{P}(A) \leq 1$, (b) $\mathrm{P}(\neg A) = 1 - \mathrm{P}(A)$, and (c) $\mathrm{P}(A \vee B) = \mathrm{P}(A) + \mathrm{P}(B) - \mathrm{P}(A \wedge B)$.

[9] Conditionalisation only covers the case when $\mathrm{P}(A) > 0$; to deal with the case when $\mathrm{P}(A) = 0$ one can use Popper-measures or let $\mathrm{P} * A$ be a non-standard measure that assigns probability 1 (or 0) to all sentences.

$\mathcal{E}_\gamma * A$, can then be defined to be the epistemic selection function that for any sentence B picks out the set of *most plausible* B-states that are also A-states (i.e. the most plausible $A \wedge B$-states):

$$(\mathcal{E}_\gamma * A)(B) = \mathcal{E}_\gamma(A \wedge B).$$

A non-conditional sentence is accepted if it is true in all the most plausible models, i.e. letting $\top$ be an arbitrary tautology ($\mathcal{E}_\gamma(\top)$ will pick out the most plausible of *all* states) and A a non-conditional sentence:

$$\mathcal{E}_\gamma \Vert\!\!\models A \text{ if and only if } v \models A, \text{ for all } v \in \mathcal{E}_\gamma(\top).$$

The epistemic consequence relation can be defined as:

> $A_1, \ldots, A_n \Vert\!\!\models B$ if and only if $\mathcal{E}_\gamma \Vert\!\!\models A_1, \ldots, \mathcal{E}_\gamma \Vert\!\!\models A_n$ implies $\mathcal{E}_\gamma \Vert\!\!\models B$, for all $\mathcal{E}_\gamma$.

These jointly ensure that Semantic Closure, Epistemic Closure and Reverse Epistemic Closure are satisfied. The model also ensures that we have a defeasible conditional that satisfies important logical properties like modus ponens and export-import.

This is not the place to discuss the relative merits of probabilistic versus qualitatitve models. Both kinds of models allow for a rich (but by no means exhaustive) representation of evidential relations and justificational structures and are able to explain intuitive judgements about indicative conditionals quite well.

6.3.3 Impossibility Results

As noted, epistemic interpretations of conditionals often treat them as *exceptional*, particularly, by not assigning truth-conditions to them. However, could it not be the case that the conditional has truth-conditions that are such that they just happen to also satisfy the Ramsey Test? For instance, couldn't the conditional have truth-conditions that, given the axioms of probability, forced the equality:

$$\mathrm{P}(A \rightarrow B) = \mathrm{P}(B \mid A)?$$

No, Lewis [20] showed (and this result has since been strenghtened in a number of ways) that a language with a conditional that embeds just like other connectives (which we would expect if the conditional had truth-conditions) cannot, on pain of triviality, satisfy this equality. As long as the standard axioms of probability are satisfied, the equality forces the probability of A to be either 0 or 1. Gärdenfors [11] has established a similar impossibility result for the qualitative case.

So to satisfy the epistemic interpretation the conditional must truly be exceptional. This conclusion is not undermined by the fact that the above equality can be non-trivially satisfied if one allows for *gappy* truth-conditions (see Sect. 6.2.2) as gappy truth-conditions are exceptional in their own right and probability measures on gappy propositions will not in general satisfy the standard axioms of probability (see, e.g. [6] for a discussion).

6.4 Concluding Remarks

Over the past century the conditional has been the locus of a rich and diverse range of philosophical debates. Considerable progress has been made in the understanding of how conditionals can be used to express connections and dependencies, as well as in the understanding of their defeasible character. To a large extent the progress can be traced back to the careful study of the formal apparatus used to represent the underlying structures. Furthermore, the insights have shed new light on issues of independent interest, a consequence of the fact that conditionals feed on structures that are of core interest in epistemology and metaphysics in general. The area of conditionals is thus one of the success stories of formal philosophy.

Many issues remain, however. Some hold that causal relations cannot be fully analyzed by counterfactuals, and so cannot be fully analyzed by similarity relations between worlds. Given that many counterfactuals are parasitic on the underlying causal structure of the world, this suggests that there may be alternative representations of this structure that could serve as the semantic basis of counterfactuals (e.g. Pearl's [27] work on causal models and Leitgeb's [17] work on probabilistic models to name just two such alternatives). Here there is a lot of work to be done. Likewise, the status of the indicative conditional under the epistemic interpretation is far from settled: Can one find representations that elegantly account for their special nature? Representations that also account for their semantic connections to counterfactuals?

References

* Indicates recommended reading.

1. Adams, E. W. (1975). *The logic of conditionals*. Dordrecht: Reidel.
2. Adams, E. W. (1998). *A primer of probability logic*. Stanford: CSLI.
3. *Arló-Costa, H. (2009). The logic of conditionals. In E. N. Zalta (Ed.), *The stanford encyclopedia of philosophy* (Spring 2009 ed.).
4. *Bennett, J. F. (2003). *A philosophical guide to conditionals*. Oxford: Oxford University press.
5. *Bradley, R. (2002). Indicative conditionals. *Erkenntnis, 56*, 345–378.
6. Bradley, R. (2007). A defence of the Ramsey test. *Mind, 116*(461), 1–21.
7. Cantwell, J. (2008). Indicative conditionals: Factual or epistemic? *Studia Logica, 88*, 157–194.
8. Cantwell, J., Lindström, S., & Rabinowicz W. (2017). McGee's counterexample to the Ramsey test. *Theoria, 83*(2), 154–168.
9. *Edgington, D. (1995). On conditionals. *Mind, 104*, 235–329.
10. Evans, J., & Over, D. (2004). *If.* Oxford: Oxford University Press.

11. Gärdenfors, P. (1986). Belief revision and the Ramsey test for conditionals. *Philosophical Review, 95*, 81–93.
12. *Grice, H. P. (1989). Indicative conditionals. In *Studies in the way of words* (pp. 58–87). Cambridge, MA: Harvard University Press.
13. Jackson, F. (1979). On assertion and indicative conditionals. *Philosophical Review, 88*, 565–589.
14. Kratzer, A. (1977). What 'must' and 'can' must and can mean. *Linguistics and Philosophy, 1*(3), 337–356.
15. Kratzer, A. (1979). Conditional necessity and possibility. In R. Bäuerle, U. Egli & A. von Stechow (Eds.), *Semantics from different points of view* (pp. 117–147). Berlin: Springer.
16. *Kratzer, A. (2012). *Modals and conditionals: New and revised perspectives*. Oxford: Oxford University Press.
17. Leitgeb, H. (2012). A probabilistic semantics for counterfactuals. *The Review of Symbolic Logic, 5*, 26–84.
18. Levi, I. (1996). *For the sake of the argument*. Cambridge: Cambridge University Press.
19. *Lewis, D. (1973). *Counterfactuals*. Cambridge: Harvard University Press.
20. *Lewis, D. (1976). Probabilities of conditionals and conditional probabilities. *Philosophical Review, 85*, 297–315.
21. Lewis, D. (1979). Counterfactual dependence and Time's arrow. *Nous, 13*, 455–476.
22. Lewis, D. (1986). Probabilities of conditionals and conditional probabilities II. *Philosophical Review, 95*, 581–589.
23. McDermott, M. (1996). On the truth conditions of certain 'If'-sentences. *Philosophical Review, 105*, 1–37.
24. *McGee, V. (1989). Conditional probabilities and compounds of conditionals. *Philosophical Review, 98*, 485–542.
25. McGee, V. (2000). To tell the truth about conditionals. *Analysis, 60*, 107–111.
26. Morton, A. (2004). Against the Ramsey test. *Analysis, 64*(4), 294–299.
27. Pearl, J. (2000). *Causality: Models, reasoning, and inference*. Cambridge: Cambridge University Press.
28. Ramsey, F. P. (1931). *Foundations of mathematics and other essays*. New York: Routledge and Kegan Paul.
29. Stalnaker, R. (1968). A theory of conditionals. In *Studies in logical theory* (No. 2 in American philosophical quarterly monograph series). Oxford: Blackwell.
30. Stalnaker, R. (1975). Indicative conditionals. *Philosophia, 5*, 269–286. Reprinted in *Ifs*, ed. W. L. Harper, R. Stalnaker & G. Pearce 1976 by Reidel.
31. Stalnaker, R. (1984). *Inquiry*. Cambridge, MA: The MIT press.
32. Strevens, M. (2003). Against Lewis's new theory of causation: A story with three morals. *Pacific Philosophical Quarterly, 84*(4), 398–412.