1.2. What is said: propositions

1.2.0. Overview

In 1.1.5, we saw the close relation between two properties of a deductive inference: (i) it is a transition from premises to conclusion that is free of any risk of new error, and (ii) the information provided by the conclusion of a deductive inference is already present in its premises. The relation between these properties points to a way of understanding the informational content of a sentence.

1.2.1. Truth values and possible worlds
First we look more closely at the concepts of risk and error involved in the idea of risk-free inference.

1.2.2. Truth conditions and propositions
We can use these ideas to give an account of the content or the meaning of a sentence, an account of what it says.

1.2.3. Ordering by content
When there is a risk-free inference from one sentence to another, the first says everything the second does, but it may say more by ruling out some possibility the second leaves open.

1.2.4. Equivalence in content
Implication in both directions between sentences shows that each says everything the other does—that is, that they say the same thing.

1.2.5. The extremes of content
Two extremes in the ordering of sentences by content are sentences that say nothing and sentences that say too much to distinguish among possibilities.

1.2.6. Logical space and the algebra of propositions
Deductive logic can be seen as the theory of the meanings of sentences in the way that arithmetic is the theory of numbers.

1.2.7. Contrasting content
Other logical relations between sentences concern differences rather than similarities in content.

1.2.8. Deductive relations in general
The relations we have considered provide a complete collection of logical relations between two sentences, and certain connections among these relations can be depicted in a traditional diagram known as the "square of opposition".

1.2.1. Truth values and possible worlds

When an inference is deductive, its conclusion cannot be in error unless there is an error somewhere in its premises. The sort of error in question lies in a statement being false, so to know that an argument is valid is to know that its conclusion must be true unless at least one premise is false. Similarly, to know that a set of sentences is inconsistent—to know that it's members are deductively incompatible—is to know that these sentences cannot all be true, and to know that a set is exhaustive is to know that its members cannot all be false. This means that the ideas of truth and falsity have a central place in deductive logic, and it will be useful to have some special vocabulary for them.

It is standard to speak of truth and falsity together as *truth values* and to abbreviate their names as **T** and **F**, respectively. This gives us a way of displaying the pattern of truth values for its premises and conclusion that validity guarantees we will not encounter; it is shown in Figure 1.2.1-1.

$$\begin{array}{rc} & \mathbf{T} \\ & \mathbf{T} \\ \textit{premises:} & \vdots \\ & \mathbf{T} \\ \hline \textit{conclusion:} & \mathbf{F} \end{array}$$

Fig. 1.2.1-1. The pattern of truth values that is not a risk when an argument is valid.

Since to speak of no risk of error is to speak of no possibility of error, it is also useful to have some vocabulary for speaking of possibility and impossibility. The sort of possibility in question in deductive logic is very weak and the corresponding sort of impossibility is very strong. We will refer to this as *logical* possibility and impossibility.

Logical possibility must be far more inclusive than the sense of the term possible in most ordinary uses. A description of a situation that runs counter to the laws of physics (for example, a locomotive floating 10 feet above the earth's surface without any abnormal forces acting on it) is naturally said to be impossible. But it need not be logically impossible, and we must consider many such *physical* impossibilities when deciding whether a conclusion is deductively valid. For, otherwise, anything following from the laws of nature, including the laws themselves, would be a valid conclusion from any premises whatsoever, and these laws would not say anything more than mere descriptions of the facts they were designed to explain. In short, if a sentence φ ever goes beyond a set of premises—if it ever provides new information—then it is logically possible for φ to be false. Or, to put it another way, a situation is logically possible if it can be coherently described. So science fiction and fantasy that is not actually self-contradictory will be logically possible.

We can say that something is impossible by saying that "there is no possibility" of it being true. In saying this, we use a form of words analogous to one we might use to say that there is no photograph of Abraham Lincoln chopping wood. That is, in saying "there is no possibility," we speak of possibilities as if they were things like photographs. This way of speaking about possibilities is convenient, so it is worth spending a moment thinking about what sort of things possibilities might be. The sort of possibility of chief interest to us is a complete state of affairs or state of the world, where this is understood to include facts concerning the full course of history, both past and future, throughout the universe. Since the late 17th century philosopher Leibniz, philosophers have used the phrase *possible world* as a particularly graphic way of referring to possibilities in this sense. For instance, Leibniz held that the goodness of God implied that the actual world must be the best of all possible worlds, and by this he meant that God made the entire course of history as good as it was logically possible for it to be.

1.2.2. Truth conditions and propositions

When judging the validity of an argument, what we need to know about its premises and conclusion are the truth values of these sentences in various possible worlds. This information about a sentence is an aspect of its meaning that we will call its *truth conditions*. That is, when we are able to tell, no matter what possible world we might be given, whether or not a sentence is true, we know the conditions under which the sentence is true; and, when we know those conditions, we can tell whether or not it is true in a given possible world.

It will also be convenient to be able to speak of this kind of meaning or aspect of meaning as an entity in its own right. We will do this by speaking of the truth conditions of a sentence as encapsulated in the *proposition* expressed by the sentence. This proposition can be thought of as a way of dividing the full range of possible worlds into those in which the sentence is true and those in which it is false. And we can picture a proposition as a division of an area representing the full range of possibilities into two regions.

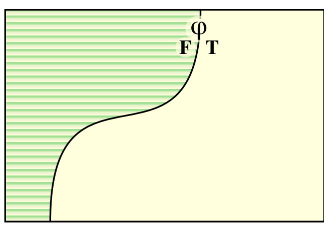

Fig. 1.2.2-1. The proposition expressed by a sentence φ, seen as dividing the full range of possible worlds into possibilities in which it is true and possibilities in which it is false.

Since knowing what possibilities are in one of these regions tells us that the rest are in the other region, we know what proposition is expressed by a sentence when we know what possibilities it rules out—or know what possibilities it leaves open. It might seem that the proposition is really indicated by the line between the two. And that's right provided we add an indication of which side of the line corresponds to truth and which side to falsity.

We will use several ways of speaking about these two regions. On the one hand, a proposition can be said to divide the possible worlds into the possibilities it *rules out* and the ones it *leaves open*. Leaving open a possibility is a failure to rule it out, and it will sometimes be useful to have a more positive way of speaking about the possibilities in which a sentence is true: we can say in such a case that the sentence covers that possibility. So a proposition can also be seen as a division of all possible worlds into possibilities *covered* and possibilites *not covered*.

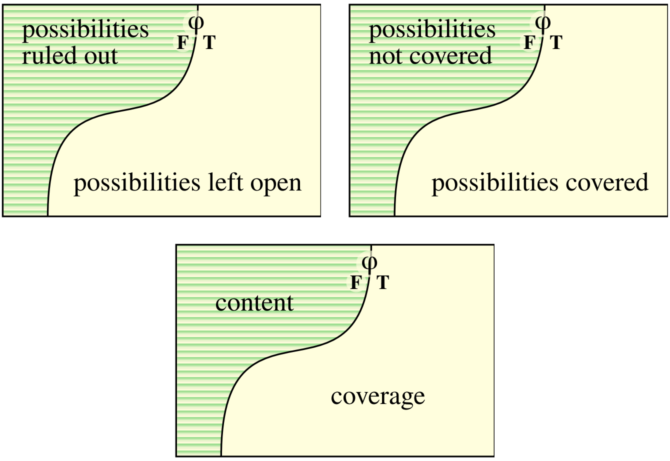

Fig. 1.2.2-2. Three ways of describing the two regions into the proposition expressed by a sentence φ divides the full range of possible worlds.

For reasons that will be discussed in the next subsection, we will speak of the collection of possibilities ruled out by a proposition as its *content*, and it is natural to refer to the full range of possibilities covered by a proposition as its *coverage*. So a proposition can be said to divide the possible worlds into its content and its coverage.

1.2.3. Ordering by content

When we judge the validity of an argument we are comparing the content of the conclusion to the contents of the premises, and the ideas of truth values and possible worlds are designed to help us speak about the basis for that comparison. We can see more of what this sort of comparison involves and what similar comparisons are possible by focusing on comparisons of two sentences.

The term *implies* is a more common English synonym of entails, and we will use it often when considering an argument that has only one premise (i.e., an "immediate inference" in traditional terminology noted in 1.1.2). Thus φ implies (or entails) ψ when there is no risk that ψ will be in error without any error in φ—i.e., when there is no logically possible world in which ψ is false even though φ is true. The impossibility of a T-F pattern of truth values—in this case for φ and ψ—is an idea that will reappear often, and we will say that a possible world that did make φ true and ψ false would have *separated* φ from ψ. So φ implies ψ when φ cannot be separated from ψ. The separation in question is separation in regards to truth: when φ implies ψ, if φ is true, then ψ will be true as well.

Separation in this sense is asymmetric. Even if φ cannot be separated from ψ, it may be possible to separate ψ from φ. For example The meeting is tomorrow morning cannot be separated from The meeting is tomorrow. But if the meeting is in fact tomorrow afternoon, the latter sentence will have been separated from the former, for it will be true that the meeting is tomorrow even though it is false that it is tomorrow morning. Clearly what is going on here is that The meeting is tomorrow morning says everything that is said by The meeting is tomorrow, and says something more. The first sentence cannot be separated from the second because, for the first sentence to be true, everything said by the second sentence must be true. But the second sentence can be separated from the first because the extra content of the first may be false even though the second sentence is true. The same point can be made in terms of coverage. If φ implies ψ, then ψ must cover all the possibilities that φ does—for otherwise φ could be separated from it—but it may cover others that φ does not. The second sentence in the example covers the possibility of an afternoon meeting but the first does not.

In short, implication is a relation of both content and coverage, but in opposite directions. If φ implies ψ, then the content of φ includes the content of ψ, and the coverage of ψ includes the coverage of φ. When the relation fails to hold in the other direction—in symbolic notation, when $\psi \nvDash \varphi$—we know that the content of φ extends beyond that of ψ. That's why there can be a possibility separating ψ from φ, a possibility where the extra content of φ is false even though the content of ψ is true. And such a possible world will be part of the coverage of ψ but not that of φ, so the coverage of ψ extends beyond that of φ.

As a more extended example of this terminology, consider the following series of successively more informative statements, each implied by the one below it:

The package will arrive sometime
⫫ (i.e., *is implied by*)
The package will arrive next week
⫫ (*is implied by*)
The package will arrive next Wednesday
⫫ (*is implied by*)
The package will arrive next Wednesday morning

Content increases as we go down the list, and coverage decreases. Each sentence above the last covers some possibilities that are ruled out by the sentence below it. And in general, as we add information, we reduce the range of possibilities that are covered. We will often speak of a sentence that rules out possibilities another does not (and thus does not cover possibilities that the other does) as making a *stronger* claim, and we will speak of sentence that does not rule out possibilities ruled out by another (and thus covers possibilities the other does not) as making a *weaker* claim. So, in the list above, the sentences closer to the bottom make the stronger claims and those closer to the top make the weaker ones.

The relation between a sentence expressing a stronger proposition and a sentence expressing a weaker can be displayed graphically by using the depiction of a proposition as a line between the possibilities it rules out and those it leaves open.

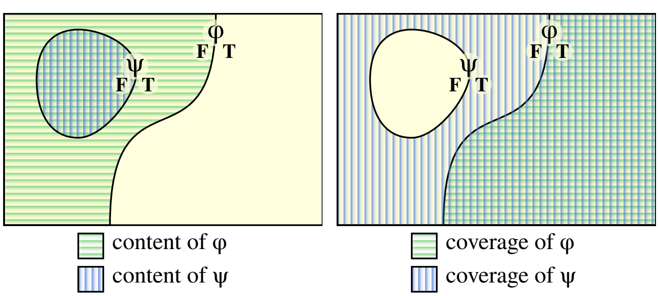

Fig. 1.2.3-1. The relation between non-equivalent propositions φ and ψ where $\varphi \vDash \psi$, depicted (on the left) by indicating the relation between the possibilities ruled out and (on the right) by indicating the possibilities left open by φ and ψ.

Here $\varphi \vDash \psi$ but $\psi \nvDash \varphi$, so there are no possible worlds where φ is true and ψ is false, but there are possibilities where φ is false while ψ is true. The possibilities ruled out by ψ, which are also ruled out by φ, are in the small region hatched in both directions on the upper left. The area outside this region but still on the left of the line running through the middle of rectangle are possibilities ruled out by φ that are left open by ψ. These possibilities separate ψ from φ and thus show that $\psi \nvDash \varphi$. The diagram on the right depicts that same relation by way of possibilities covered rather than possibilities ruled out. While φ covers only those possibilities on the right of the diagram, ψ covers all that are not within the region at the upper left, so ψ covers any possibility that φ does but not vice versa; that is, ψ is true whenever φ—i.e., φ cannot be separated from it—but ψ is true in some possible worlds in which φ is false—i.e., it can be separated from φ.

If the relation of implication holds in both directions—if both $\varphi \vDash \psi$ and $\psi \vDash \varphi$—then each of the two sentences says everything the other does, so they provide exactly the same information and cover the same possibilities, differing at most in their wording. For example, although one of the sentences Sam lives somewhere in northern Illinois or southern Wisconsin and Sam lives somewhere in southern Wisconsin or northern Illinois might be chosen over the other depending on the circumstances, they allow the same possibilities for Sam's residence and thus provide the same information about it. We will say that sentences that have the same informational content are *(logically) equivalent* (usually dropping the qualification logically since we will not be considering other sorts of equivalence). Our notation for logical equivalence—the sign $\simeq$ (*tilde equal*)—gets used for many different kinds of equivalence, but we will use it only for logical equivalence. The idea of logical equivalence can also be described directly in terms of truth values and possible worlds. When $\varphi \simeq \psi$, we know that neither can be separated from the other, so φ and ψ must have the same truth value as each other in any possible world. And that means that equivalent sentences have the same truth conditions and express the same proposition.

Since relations of entailment depend only on possibilities of truth and falsity, equivalent sentences entail and are entailed by the same sentences. That means that entailment can be thought of as a relation between the propositions they express. It provides a sort of ordering of propositions by their content that can be compared to the ordering of numbers by $\leq$ and $\geq$. Whether entailment seems more like $\leq$ or $\geq$ depends on whether we think of it as a comparison of possibilities left open or of possibilities ruled out. When a choice needs to be made, we'll general adopt the former perspective. In any case, the analogy is with $\leq$ or $\geq$ rather than $<$ or $>$ because $\varphi \vDash \psi$ tells us that φ says more *or the same as* ψ, that it leaves fewer *or the same* possibilities open.

We have been employing analogies between implication and numerical ordering and the related sorts of comparison that are associated with terms like stronger and weaker. These analogies rest on two properties that implication shares with many other relations. First of all, it is *transitive* in the sense that implication by a premise φ carries over from a valid conclusion ψ to any sentence χ implied by that conclusion: if $\varphi \vDash \psi$ and $\psi \vDash \chi$, then $\varphi \vDash \chi$. That is, we do not count steps in a chain of related items (as is done with parent of, grandparent of, etc., which are not transitive relations) but simply report the existence of a chain no matter what its length (as is done with ancestor of, which is transitive).

Just about any relation that we would be ready to call an "ordering" is transitive. Implication also shares with certain orderings the more special property of being *reflexive* in the sense that every sentence implies itself. Reflexivity is what distinguishes orderings like $\leq$ and as strong as or stronger than from $<$ and stronger than. In the first two, examples reflexivity is achieved by tacking on a second reflexive relation ($=$ in one case and equally strong as in the other) as an alternative. The analogous relation in the case of implication (i.e., one amounting to "equal in content to") is equivalence, but that is an alternative already built into implication (i.e., one sort of case in which a sentence φ implies a sentence ψ is when they are equivalent), so it does not need to be added.

Relations like equality ($=$), the relation equally strong as, and the relation of logical equivalence are reflexive and transitive, but they are not very effective in ordering things because they have no direction: if they hold between a pair of things in one direction, they hold in the other direction, too. In particular, if $\varphi \simeq \psi$ then $\psi \simeq \varphi$. A relation with this property is said to be *symmetric*. Relations with the three properties of transitivity, reflexivity, and symmetry are said to be *equivalence relations*. Any equivalence relation points to equivalence or equality in some respect, and different equivalence relations point to different sorts of equality or equivalence. Logical equivalence between sentences points to equivalence in content or in the proposition expressed.

1.2.4. Equivalence in content

If the relation of implication holds in both directions—if both $\varphi \vDash \psi$ and $\psi \vDash \varphi$—then each of the two sentences says everything the other does, so they provide exactly the same information and cover the same possibilities, differing at most in their wording. For example, although one of the sentences Sam lives somewhere in northern Illinois or southern Wisconsin and Sam lives somewhere in southern Wisconsin or northern Illinois might be chosen over the other depending on the circumstances, they allow the same possibilities for Sam's residence and thus provide the same information about it. We will say that sentences that have the same informational content are *(logically) equivalent* (usually dropping the qualification logically since we will not be considering other sorts of equivalence). Our notation for logical equivalence—the sign $\simeq$ (*tilde equal*)—gets used for many different kinds of equivalence, but we will use it only for logical equivalence. The idea of logical equivalence can also be described directly in terms of truth values and possible worlds. When $\varphi \simeq \psi$, we know that neither can be separated from the other, so φ and ψ must have the same truth value as each other in any possible world. And that means that equivalent sentences have the same truth conditions and express the same proposition.

Since relations of entailment depend only on possibilities of truth and falsity, equivalent sentences entail and are entailed by the same sentences. That means that entailment can be thought of as a relation between the propositions they express. It provides a sort of ordering of propositions by their content that can be compared to the ordering of numbers by $\leq$ and $\geq$. Whether entailment seems more like $\leq$ or $\geq$ depends on whether we think of it as a comparison of possibilities left open or of possibilities ruled out. When a choice needs to be made, we'll general adopt the former perspective. In any case, the analogy is with $\leq$ or $\geq$ rather than $<$ or $>$ because $\varphi \vDash \psi$ tells us that φ says more *or the same as* ψ—or, in other words, that it covers fewer *or the same* possibilities.

We have been employing analogies between implication and numerical ordering and the related sorts of comparison that are associated with terms like stronger and weaker. These analogies rest on two properties that implication shares with many other relations. First of all, it is *transitive* in the sense that implication by a premise φ carries over from a valid conclusion ψ to any sentence χ implied by that conclusion: if $\varphi \vDash \psi$ and $\psi \vDash \chi$, then $\varphi \vDash \chi$. That is, we do not count steps in a chain of related items (as is done with parent of, grandparent of, etc., which are not transitive relations) but simply report the existence of a chain no matter what its length (as is done with ancestor of, which is transitive).

Just about any relation that we would be ready to call an "ordering" is transitive. Implication also shares with certain orderings the more special property of being *reflexive* in the sense that every sentence implies itself. Reflexivity is what distinguishes orderings like $\leq$ and as strong as or stronger than from $<$ and stronger than. In the first two, examples reflexivity is achieved by tacking on a second reflexive relation (= in one case and equally strong as in the other) as an alternative. The analogous relation in the case of implication (i.e., one amounting to "equal in content to") is equivalence, but that is an alternative already built into implication (i.e., one sort of case in which a sentence φ implies a sentence ψ is when they are equivalent), so it does not need to be added.

Relations like equality (=), the relation equally strong as, and the relation of logical equivalence are reflexive and transitive, but they are not very effective in ordering things because they have no direction: if they hold between a pair of things in one direction, they hold in the other direction, too. In particular, if $\varphi \simeq \psi$ then $\psi \simeq \varphi$. A relation with this property is said to be *symmetric*. Relations with the three properties of transitivity, reflexivity, and symmetry are said to be *equivalence relations*. Any equivalence relation points to equivalence or equality in some respect, and different equivalence relations point to different sorts of equality or equivalence. Logical equivalence between sentences points to equivalence in content or in the proposition expressed.

1.2.5. The extremes of content

There are two extreme examples of truth conditions or propositions. A sentence that is true in all possible worlds says nothing. It has no informational content because it leaves open all possibilities and rules nothing out. For example, the weather "forecast" Either it will rain or it won't has no chance of being wrong and is, therefore, completely worthless as a prediction. We will say that such a sentence is a *tautology*. Although there are many (indeed, infinitely many) tautologies, all express the same proposition; and the words that they use to express it are beside the point since they all say nothing in the end. In short, any two tautologies are logically equivalent. It will be convenient to establish a particular tautology and mark it by special notation. We will call this sentence *Tautology* and use the sign $\top$ (*down tack*) as our notation for it. Since the logical properties and relations we will consider depend only on the propositions expressed by sentences, any logical property or relation of $\top$ will hold for all tautologies, and we will often simply speak of $\top$ in order to say things about tautologies generally.

At the other extreme of truth conditions from tautologies are sentences that rule out all possibilities. The fact that such a sentence is the opposite of a tautology might suggest that it is maximally informative, but it sets an upper bound on informativeness in a different way: any genuinely informative sentence must say less than it does. The ultimate aim of providing information is to narrow down possibilities until a single one remains, for this would provide a complete description of the history of the universe. To go beyond this would leave us with nothing because there is no way to distinguish among possibilities if all are ruled out. For example, the forecast It will rain, but it won't is far from non-committal since it stands no chance of being right, but it is no more helpful than a tautologous one.

Sentences that rule out all possibilities make logically impossible claims, and we will refer to them as *absurd*. As was the case with tautologies, any two absurd sentences are logically equivalent. So, as with tautologies, we will introduce a particular example of an absurdity, named *Absurdity*, and we use the special notation $\bot$ (the perpendicular sign, or *up tack*) for it.

A tautology is implied by any sentence φ since, as it rules out no possibilities, it must cover any possibility that is covered by φ. The sentence $\top$ is thus the weakest sentence there could be and it can stand at the top of any ordering by logical strength like that depicted in 1.2.3 . Analogously, an absurd sentence implies all sentences: since it covers no possibilities, its coverage is included in that of any other sentence. So the sentence $\bot$ can stand at the bottom of any ordering by logical strength.

Of course, most sentences are neither tautologies nor absurd. We will say of sentence that is neither that it is *logically contingent* because there is at least one possible world in which it is true and at least one where it is false, so its truth or falsity is not settled by logic.

1.2.6. Logical space and the algebra of propositions

Logic is concerned with propositions in the way mathematics is concerned with numbers, but propositions are not numbers. While numbers can be ordered in a linear way, the collection of propositions has a more complex structure. The series of examples of increasing strength we looked at in 1.2.3 did form a single chain, but it should be clear that we could have gone in many different directions to find stronger or weaker claims propositions. For example, The package will arrive next Wednesday is implied by The package will arrive next Wednesday morning but also by The package will arrive next Wednesday afternoon, and neither of the latter sentences implies the other. And The package will arrive next Wednesday implies the sentences The package will arrive next week and The package will arrive on a Wednesday, and the latter two sentences are not ordered one way or the other by implication.

This metaphor of many directions suggests a space of more than one dimension; and, although the structure of a collection of propositions differs not only from the 1-dimensional number line but also from the structure of ordinary 2- or 3-dimensional space, spatial metaphors and diagrams can help in thinking about its structure. These metaphors and can be associated with the term *logical space* that was introduced by the philosopher Ludwig Wittgenstein (1889-1951).

We will actually use two different sorts of spatial metaphor. One metaphor is the one used in 1.2.2 to depict propositions. In it, possible worlds are the points of logical space, and propositions determine regions in the space by drawing a boundary between the possibilities they rule out and the ones they leave open. But we use a different metaphor when we speak of increasing strength in many different directions. According to this second metaphor, propositions are points in space rather than regions, and possible worlds function in it behind the scenes as something like the dimensions of the space. If we were to apply this idea in any very realistic way, the space would have too many dimensions to be visualized, but in artificially simple cases this sort of space can be depicted by a figure in ordinary 2- or 3-dimensional space.

Let's begin to look further at these ideas by considering an very simple example of the first sort of logical space. Suppose there were only 4 possible worlds. A proposition will either rule out or leave open each of these possibilities. Figure 1.2.6-1 is intended to illustrate two such propositions. Each of these propositions rules out two of the four possibilities (in the hatched areas) and leaves open two others. The proposition expressed by the sentence φ rules out the two possibilities at the bottom of the diagram and the one expressed by ψ rules out the ones at the right. As a result both rule out the possible world in the lower right of the diagram and neither rules out the one in the upper left.

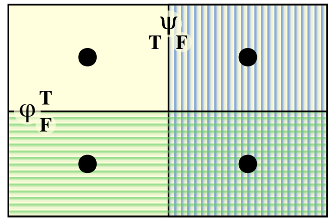

Fig. 1.2.6-1. The possibilities (the hatched bottom and right halves) that are ruled out by two propositions.

Of course, these are not the only propositions that can be expressed given this range of possibilities. A proposition has two options for each possible world: it may rule it out or leave it open. With 4 possible worlds this means that there are $2 \times 2 \times 2 \times 2 = 16$ propositions in all, and 6 of these rule out exactly two possible worlds.

We can illustrate all 16 of these propositions by using a logical space of the second sort. Figure 1.2.6-2 depicts (in two dimensions) a 3-dimensional figure that is one possible representation of a 4-dimensional cube. It is labeled to suggest what sorts of sentences might express these propositions.

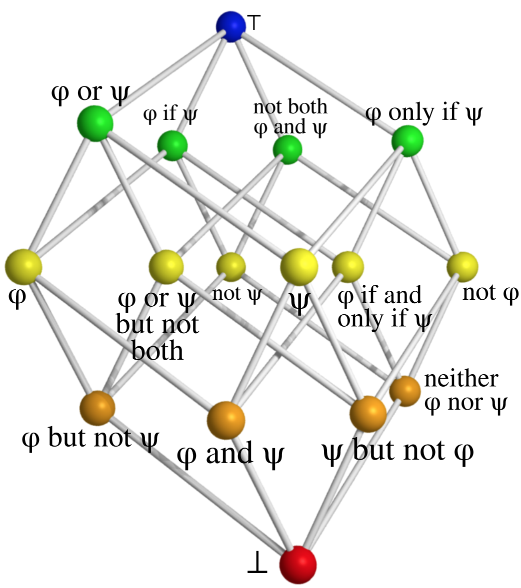

Fig. 1.2.6-2. The sixteen propositions when there are 4 possible worlds.

You can imagine that the propositions φ (which appears at the left) and ψ (near the center) are the two propositions depicted in Figure 1.2.6-1.

The levels in the structure correspond to grades of strength, with Absurdity at the bottom ruling out all possible worlds and Tautology at the top ruling out none. A line connects propositions that differ only with respect to one possible world. This world separates the proposition higher in the diagram from the one below it, but the lower proposition implies the one above it. Each of the four propositions immediately above Absurdity then leaves open just one possible world. Lines connecting propositions that differ with respect to this world are parallel (in the 3-dimensional figure, though not in the 2-dimensional perspective image on the screen or page) to the line connecting the proposition to Absurdity. In this sense, the worlds can be thought of as the dimensions on which the content of propositions can vary.

The relation between the two sorts of diagram can be seen by replacing each proposition in Figure 1.2.6-2 by its representation using a diagram of the sort illustrated in Figure 1.2.6-1. Putting the two sorts of illustration together in this way gives us the following picture of the same 16 propositions.

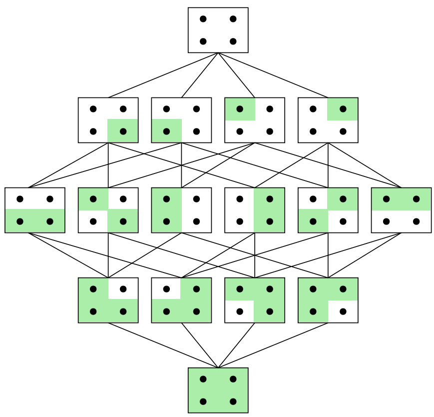

Fig. 1.2.6-3. The propositions generated by 4 possible worlds depicted as regions in one logical space (the repeated rectangle) and as points in another (the overall diagram).

The spacing of the nodes differs between Figures 1.2.6-2 and 1.2.6-3 but the left-to-right order at each level is the same, and the regions associated with φ and ψ are the same as those depicted in Figure 1.2.6-1. Since a sentence that rules out more possibilities makes a stronger claim, the size of the region occupied by the possibilities it rules out can be thought to correspond to the strength of the claim it makes. Notice that the regions ruled out here increase towards the bottom of the diagram and that they are the same in size for all nodes on the same level.

The whole structure of Figure 1.2.6-2 can be seen as a complex diamond formed of four diamonds whose corresponding vertices are linked. A simple diamond is the structure of the $2 \times 2 = 4$ propositions we would have with only

2 possible worlds. The structure in Figure 1.2.6-2 doubles the number of possible worlds and squares the number of propositions. If we were to double the number of possible worlds again to 8, we would square the number of propositions to get 256. The structure they would form could be obtained by replacing each node in the structure of Figure 1.2.6-2 by a small structure of the same form and replacing each line by a bundle of 16 lines connecting the corresponding nodes.

To get a sense of the structure of the set of propositions for a realistically large set of possible worlds, imagine carrying out this process over and over again. The result will always have an upper and lower limit (⊤ and ⊥) and many different nodes on each of its intermediate levels. As the number of possible worlds increases, the distribution of possible worlds among the various degrees of strength (which is 1, 4, 6, 4, 1 in Figure 1.2.6-2) will more and more closely approximate a bell curve. But the bell shape of the curve will also narrow significantly, and bulk of the propositions will be found in intermediate degrees of strength. In short, as the space of propositions gets closer to a realistic degree of complexity, it departs further and further from a single line with ⊤ at the top and ⊥ at the bottom.

1.2.7. Contrasting content

We arrived at the relation of implication by considering entailment by a single premise. If we do the same with exclusion, we arrive at another relation between sentences. If φ excludes ψ, then the set {φ, ψ} formed of the two is inconsistent. When sentences φ and ψ are related in this way, it is equally true that ψ excludes φ. This reversability of this relation is reflected in the usual terminology for it: when there is no possible world in which φ and ψ are together true, φ and ψ are said to be *mutually exclusive*. There is no standard notation for the relation, and we will eventually have a way of expressing it in terms of entailment; but, when it is convenient to have special notation, we will write φ △ ψ to say that φ and ψ are mutually exclusive. This use of the *up-pointing triangle* is intended simply to reflect the shape of signs for some related ideas. One of these related ideas is Absurdity. When the possibilities ruled out by a pair of mutually exclusive sentences are taken together—when their contents are added up—they include all possibilities whatsoever. In this respect, mutually exclusive sentences together do what ⊥ does on its own.

Mutually exclusive sentences provide one example of the differences in propositions that made for the horizontal spread of the logical space of Figure 1.2.6-2. Indeed, one of the examples cited there, the sentences The package will arrive next Wednesday morning and The package will arrive next Wednesday afternoon was a pair of mutually exclusive sentences. Mutually exclusive sentences differ to the extent that there is no overlap in their coverage (since they cannot be both true in any possible world). From one point of view, that is a pretty considerable difference; but, as this example illustrates, such sentences can still have a lot in common. And, in general, sentences that rule out many possibilities may express propositions that divide the space of possibilities in very similar ways even though they have no overlap in their coverage. (Imagine the whole content of an encyclopedia bundled up in a single proposition; and then imagine two encyclopedias that differ only in their reports of the population of a single town—and perhaps differ only by one person in their reports of this population.)

The diagrams below depict mutually exclusive sentences φ and ψ. Notice that in none of the three regions shown are both true. The diagram on the left shows that, when the contents of the two sentences are added up, all possibilities are included. In this sense, the relation of mutual exlcusivity is an indication of the strength of the two taken together: any possibility not ruled out by one is ruled out by the other. The diagram on the right compares the coverage of the two, and we see the lack of overlap.

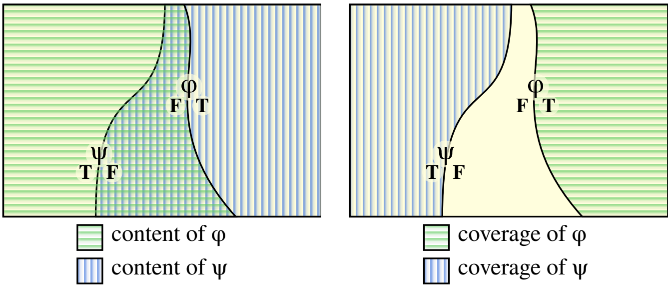

Fig. 1.2.7-1. The relation betwen mutually exclusive sentences φ and ψ, depicted on the left in terms of the possibilities each rules out and, on the right, in terms of the possibilities each covers.

The region in the middle of the diagram could be contracted to a line and the sentences would still be mutually exclusive, for then the sentences would still combine to rule out all possibilites and would still show no overlap in coverage.

This suggests a distinction that may be made among pairs of mutually exclusive sentences. All mutually exclusive sentences are opposed to one another, and they can be thought of as opposites. But there are different sorts of opposites. Some, like The glass is full and The glass is empty are extremes that may both fail in intermediate cases, and the example depicted above is like this. Others, like The glass is full and The glass is not full cover all the ground between them and do not leave room for a third alternative. Opposites of the latter sort might be described as *exactly* opposite.

The difference between these sorts of opposition is tied to another way in which sentences can differ. Sentences φ and ψ are *jointly exhaustive*—that is, together they form an exhaustive set—when there is no possible world in which both are false, when there is no possible world that both rule out. If we put together the possibilities covered by such sentences, the result will include all possibilities because any possibility not covered by one must be covered by the other; and, in this sense, these sentences jointly exhaust all possibilities. Such sentences certainly differ in meaning—since there is no overlap in their content, they can be said to have no content in common—but they are not opposites in the sense of being incompatible. They might be thought of instead as *complementary* since, in regard to coverage, each picks up where the other leaves off. We will use a *down-pointing triangle* ▽ as our notation for this relation, as in the case of △ because of the similarity in shape between ▽ and some ideas related to joint exhaustiveness. Tautology is one of these ideas: in regard to coverage, jointly exhaustive sentences do together what ⊤ does on its own.

In the diagrams below, the absence of common content is depicted on the left. On the right, we see how the areas of coverage of the two sentences combine to exhaust all possibilities whatsoever.

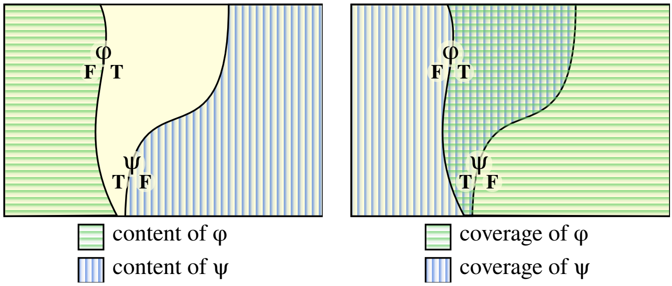

Fig. 1.2.7-2. The relation betwen jointly exhaustive sentences φ and ψ, depicted on the left in terms of the possibilities each rules out and, on the right, in terms of the possibilities each covers.

As with the depiction of mutually exclusive sentences, the region in the middle could be contracted to a line.

This situation is depicted in the following diagram.

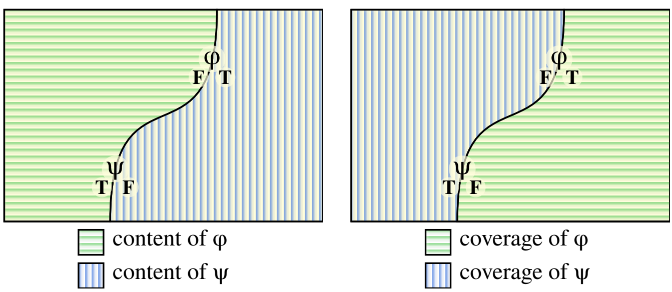

Fig. 1.2.7-3. The relation betwen sentences φ and ψ which are contradictory —i.e., both mutually exclusive and jointly exhaustive—depicted on the left in terms of the possibilities each rules out and, on the right, in terms of the possibilities each covers.

When sentences are not only mutually exclusive but also jointly exhaustive, we will say that they are *contradictory*. Contradictory sentences—like The glass is full and The glass is not full—are bound to have opposite truth values. We will write φ ⋈ ψ to say that φ and ψ are contradictory (using the symbol *bowtie*). (You might think of the symbol as indicating that things get turned upside down when moving from one sentence to the other.)

Although our use of the term contradictory is the standard one in discussions of deductive logic, in ordinary speech this term is often applied to sentences that are only mutually exclusive. In particular, when a claim is said to be "self-contradictory," what is meant is that part of what it says excludes something else it says. Such a sentence will not contradict itself in the sense in which we will use the term because that would require that it be both true and false in each possible world, and that cannot happen if there are any possible worlds at all (an assumption we can feel safe in making).

Just as the propositions expressed by logically strong sentences need not be far different even when they are mutually exclusive, the propositions expressed by logically weak sentences need not be far different even when they are jointly exhaustive. It is contradictory sentences that provide the true extreme examples of difference. When logical space in Figure 1.2.6-2 is thought of in three dimensions, the contradictory sentences appear in diametically opposite positions. Notice that mutually exclusive sentences cannot both appear above the middle level (sentences above the middle cover more than half the possibilities, so any two must have overlapping coverage), and jointly exhaustive sentences cannot appear both below the middle. Contradictory sentences fall under both restrictions. A pair of contradictory sentences might both appear on the middle level, but it is also possible for one to be of more than average logical strength if the other is relatively weak. The extreme case of this is provided by $\bot$ and $\top$, which are contradictory.

1.2.8. Deductive relations in general

The six basic deductive relations between two sentences that we have considered are shown in the following table:

relation	*pattern ruled out*		*relation*
$\varphi \vDash \psi$	φ is **T**	ψ is **F**	$\varphi \simeq \psi$
$\varphi \Dashv \psi$	φ is **F**	ψ is **T**	
$\varphi \vartriangle \psi$	φ is **T**	ψ is **T**	$\varphi \bowtie \psi$
$\varphi \triangledown \psi$	φ is **F**	ψ is **F**	

Each says that one or more patterns of truth values occurs in no possible worlds. And there are no other ways of doing this that yield genuine relations between a pair of sentences. If we rule any pair of patterns other the pairs ruled out by equivalence and contradictoriness, we end up specifying the truth value of one of the two sentences—i.e., we say of either φ or ψ that is a tautology or that it is absurd. And any way of ruling out three patterns must do this for both φ and ψ.

When no deductive relation holds between a pair of sentences φ and ψ—that is, when each of four patterns of truth values for the two appears in some possible world—we will say that φ and ψ are *logically independent*. Not only are logically independent sentences unordered by implication, they are not mutually exclusive or jointly exhaustive. And it follows from this, of course, that they are not equivalent or contradictory and also that neither is a tautology or absurd (so each one is logically contingent). This sort of thing is true for most pairs of sentences. Although sentences on different topics almost always provide examples, logically independent sentences do not need to differ in subject matter. For example, the sentences The package will arrive next week and The package will arrive on a Wednesday (a pair of sentences mentioned in 1.2.6) are logically independent since it is possible for the package to arrive next week but not on Wednesday, for it to arrive on a Wednesday but not next week, for it to arrive next Wednesday, and for it to arrive neither next week nor on a Wednesday.

There are a number of connections among the six deductive relations that can be depicted in a traditional form of diagram known as a *square of opposition*. In the case of the examples that were used to illustrate various sorts of opposites, the square can be arranged as shown in Figure 1.2.8-1. The vertical structure of the diagram displays

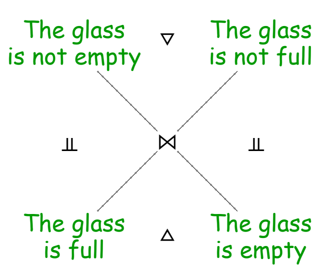

Fig. 1.2.8-1. A square of opposition.

ordering by implication in the way we have before: each of the sentences in the bottom row implies the sentence show above it. The horizontal structure of the diagram displays the sorts of opposition. The sentences along the bottom are mutually exclusive, those along the top are jointly exhaustive, and the sentences along the diagonals are contradictory.

Given one side of the square, the other side can be reconstructed by taking contradictories. For example if $\varphi \vDash \psi$, then φ will be mutually exclusive with any sentence contradictory to ψ and any sentence contradictory to φ will be jointly exhaustive with ψ. This provides a way of generating squares of opposition, but it also shows something more important: implication and contradictoriness can be seen as the fundamental deductive relations between pairs of sentences. There is more to be said about deductive relations when we consider larger collections of sentences, but we will see in 1.4.6 that something analogous continues to be true.

1.2.s. Summary

1 The relation of entailment concerns the possibilities of truth and falsity for premises and conclusions; that is, it concerns the truth values of these sentences in various possible worlds. The possibilities in question are logical possibilities, which may be understood as the situations whose description is permitted by the semantic rules of the language.

2 The deductive relations a sentence stands in depend on its truth values in various possible worlds. That is, they depend on its truth conditions. These truth conditions are encapsulated in the proposition it expresses, which can be thought of as a way of dividing all possibilities into those it rules out and those it leaves open. This means that a proposition can be depicted as a division of space into two regions.

3 Entailment by a single premise, or implication, is a relation between sentences that orders them by their content. More precisely, $\varphi \vDash \psi$ when φ says everything that is said by ψ. When this relation does not hold, it is possible for φ to remain true when something said by ψ is false; such a possibility is said to separate φ from ψ. When $\varphi \vDash \psi$ but not vice versa, φ says more than ψ and we will often say that φ makes a stronger claim and ψ a weaker one.

4 When sentences imply each other, they say the some thing, and we say they are equivalent, a relation for which we use the sign $\simeq$.

5 At one extreme are tautologies, which rule out no possibilities and thus have no content. All tautologies are equivalent and we will distinguish one, Tautology, for which we use the notation $\top$. At the other extreme are sentences that rule out all possibilities. Such sentences are absurd and all are equivalent to the single representative Absurdity, for which we use the notation $\bot$. A sentence at neither of these extremes is logically contingent.

6 Although certain groups of sentences can be ordered linearly between $\bot$ and $\top$ as a series of claims with steadily increasing content, the full range of propositions expressed by sentences are better thought of as inhabiting a much more complex logical space. This space might be a space of possibilities with propositions appearing as ways of dividing the space into regions, or it might be a space that has as its points propositions themselves. Logical space in this second sense has a bottom in the proposition expressed by $\bot$ and a top provided by $\top$. When there are a significant number of possible worlds, there will be many more propositions with intermediate content than there are strong propositions near $\bot$ or weak ones near $\top$.

7 Sentences can also be compared by describing differences in what they say.

Sentences that cannot both be true are mutually exclusive (a relation for which we use the sign $\vartriangle$). The claims made by such sentences are opposite but opposite in a way that permits a third alternative. Sentences which are complementary in the sense that each must be true if the other is false are jointly exhaustive (for which our notation is $\triangledown$). When these two relations both hold, sentences are contradictory (a relation for which we use the sign $\bowtie$). Contradictory sentences always have opposite truth values and thus make claims that are opposite in a way that permits no third alternative.

8 The relations of entailment, mutual exclusiveness, and joint exhaustiveness along with the properties of tautologousness and absurdity enable us to describe any deductive property or relation of two sentences. There are connections among entailment, mutual exclusiveness, and joint exhaustiveness that can be displayed by a square of opposition. Sentences that are neither mutually exclusive nor jointly exhaustive and neither or which implies the other are logically independent.

1.2.x. Exercise questions

1. Each of the following claims that a deductive relation holds between a pair of sentences. In each case, judge whether the claim is true and, if not, describe a sort of possibility that shows it is not true. Briefly explain your answers. For example, we can say that The package will arrive sometime does not entail The package will arrive next week because the possibility that it will arrive before or after next week is ruled out by the conclusion but not by the premise. In answering, it is safe to understand the sentences below all in the most straightforward way; you will miss the point of the exercise if you try to look for subtle or obscure possibilities.

- **a.** The package will arrive next Tueday $\vDash$ The package will arrive next week
- **b.** The package will arrive next week $\vDash$ The package will arrive next Tuesday
- **c.** The package will arrive next Tueday $\vartriangle$ The package will arrive next week
- **d.** The package will arrive next Tuesday $\vartriangle$ The package will arrive next Wednesday
- **e.** The package will arrive before next Tueday $\triangledown$ The package will arrive after next Tuesday
- **f.** The package will arrive next Tuesday or before $\triangledown$ The package will not arrive before next Wednesday
- **g.** The package will arrive after next Tuesday $\simeq$ The package will arrive next Wednesday or later
- **h.** The bridge will open at the end of May $\simeq$ The bridge will open before June
- **i.** The package will arrive before next Wednesday $\bowtie$ The package will arrive after next Wednesday
- **j.** The bridge will open before June $\bowtie$ The bridge will open in June or later or never at all

2. Some of the following claims about deductive relations hold for any sentence φ, some for no sentence φ, and others hold only if φ is a tautology or only if it is absurd. In each case, say which is so and explain your answer.

a.	$\varphi \vDash \varphi$	**b.**	$\varphi \vDash \top$	**c.**	$\varphi \vDash \bot$
		d.	$\top \vDash \varphi$	**e.**	$\bot \vDash \varphi$
f.	$\varphi \triangledown \varphi$	**g.**	$\varphi \triangledown \top$	**h.**	$\varphi \triangledown \bot$

i.	$\varphi \vartriangle \varphi$	**j.**	$\varphi \vartriangle \top$	**k.**	$\varphi \vartriangle \bot$
l.	$\varphi \simeq \varphi$	**m.**	$\varphi \simeq \top$	**n.**	$\varphi \simeq \bot$
o.	$\varphi \bowtie \varphi$	**p.**	$\varphi \bowtie \top$	**q.**	$\varphi \bowtie \bot$

3. The headings at the left of the table give information about the relation of φ and ψ and those at the top give information about the relation of ψ and χ. Fill in cells of the table by indicating what, if anything, you can conclude in each case about the relation of φ and χ. For example, if $\varphi \vDash \psi$ and $\psi \vDash \chi$, we cannot have φ true and χ false, so $\varphi \vDash \chi$ (this is the transitivity of implication). However, no other patterns for φ and χ are ruled out, so "$\varphi \vDash \chi$" is the most we can say on the basis of the given information, and it can be entered in the upper left cell.

	$\psi \vDash \chi$	$\chi \vDash \psi$	$\psi \simeq \chi$	$\psi \vartriangle \chi$	$\psi \triangledown \chi$	$\psi \bowtie \chi$
$\varphi \vDash \psi$						
$\psi \vDash \varphi$						
$\varphi \simeq \psi$						
$\varphi \vartriangle \psi$						
$\varphi \triangledown \psi$						
$\varphi \bowtie \psi$						

1.2.xa. Exercise answers

1. a. The package will arrive next Tueday entails The package will arrive next week because the package arriving next Tuesday is one of ways for it to be true that it arrives next week

   b. The package will arrive next week does not entail The package will arrive next Tuesday because the premise would still be true if it arrived another day next week

   c. The package will arrive next Tuesday and The package will arrive next week are not mutually exclusive because both will be true if it does arrive next Tuesday

   d. The package will arrive next Tuesday and The package will arrive next Wednesday are mutually exclusive since the package cannot arrive both days

   e. The package will arrive before next Tueday and The package will arrive after next Tuesday are not jointly exhaustive since both will be false if it arrives on next Tuesday

   f. The package will arrive next Tuesday or before and The package will not arrive before next Wednesday are jointly exhaustive because, if the second is false—i.e., if it does arrive before next Wednesday—then the first must be true

   g. The package will arrive after next Tuesday is equivalent to The package will arrive next Wednesday or later because arriving next Wednesday or later than that are the two ways in which a package could arrive after next Tuesday

   h. The bridge will open at the end of May is not equivalent to The bridge will open before June since it is not now the end of May so the bridge could open before June by opening even earlier than the end of May

   i. The package will arrive before next Wednesday and The package will arrive after next Wednesday are not contradictory because both will be false if it arrives on next Wednesday

   j. The bridge will open before June and The bridge will open in June or later or never at all are contradictory because opening before June, opening in June, opening later than June, and not opening at all exhaust all possibilities and are mutually incompatible

2. a. $\varphi \vDash \varphi$ holds always because φ cannot fail to be true if it is true

b. $\varphi \vDash \top$ holds always because $\top$ cannot fail to be true no matter what φ is like

c. $\varphi \vDash \bot$ holds only when φ is absurd because, if there is any possibility of φ being true, there is a possibility of $\bot$ being false when φ is true

d. $\top \vDash \varphi$ holds only when φ is a tautology because if there is any possibility of φ being false, there is a possibility of it being false when $\top$ is true

e. $\bot \vDash \varphi$ holds always because there is no possibility of $\bot$ being true so no possibility of φ being false when $\bot$ is true

f. $\varphi \triangledown \varphi$ holds only when φ is a tautology because if there is any possibility of φ being false, it does not, together with itself exhaust all possibilities

g. $\varphi \triangledown \top$ holds always becuase $\top$ covers all possibilities by itself, so it certainly exhausts them when taken together with φ

h. $\varphi \triangledown \bot$ holds only when φ is a tautology becuase, since $\bot$ leaves open no possibilities, it contributes nothing to exhausting them all and φ must do that all by itself

i. $\varphi \vartriangle \varphi$ holds only when φ is absurd because, unless φ rules out all possibilities, there will be a possibility of it being true along with itself

j. $\varphi \vartriangle \top$ holds only when φ is absurd because, since $\top$ is bound to be true, any possibility of φ being true will be a possibility of both being true

k. $\varphi \vartriangle \bot$ holds always because, since $\bot$ cannot be true, it cannot be true together with any sentence (even itself)

l. $\varphi \simeq \varphi$ holds always since a sentence must have the same truth value as itself

m. $\varphi \simeq \top$ holds only when φ is a tautology because, if φ is bound to have the same truth value as a tautology, it must be one

n. $\varphi \simeq \bot$ holds only when φ is absurd because, if φ is bound to have the same truth value as an absurd sentence, it must be one

o. $\varphi \bowtie \varphi$ never holds because no sentence can be both true and false at the same time

p. $\varphi \bowtie \top$ holds only when φ is absurd because φ is bound to be false if its value is opposite that of a sentence that is bound to be true

q. $\varphi \bowtie \bot$ holds only when φ is a tautology because φ is bound to be true if its value is opposite that of a sentence that is bound to be false

3. The appearance of "—" in a cell in the table below indicates that nothing can be concluded in general about the relation between φ and χ.

	$\psi \vDash \chi$	$\chi \vDash \psi$	$\psi \simeq \chi$	$\psi \vartriangle \chi$	$\psi \triangledown \chi$	$\psi \bowtie \chi$
$\varphi \vDash \psi$	$\varphi \vDash \chi$	—†	$\varphi \vDash \chi$	$\varphi \vartriangle \chi$	—†	$\varphi \vartriangle \chi$
$\psi \vDash \varphi$	—*	$\chi \vDash \varphi$	$\chi \vDash \varphi$	—*	$\varphi \triangledown \chi$	$\varphi \triangledown \chi$
$\varphi \simeq \psi$	$\varphi \vDash \chi$	$\chi \vDash \varphi$	$\varphi \simeq \chi$	$\varphi \vartriangle \chi$	$\varphi \triangledown \chi$	$\varphi \bowtie \chi$
$\varphi \vartriangle \psi$	—*	$\varphi \vartriangle \chi$	$\varphi \vartriangle \chi$	—*	$\varphi \vDash \chi$	$\varphi \vDash \chi$
$\varphi \triangledown \psi$	$\varphi \triangledown \chi$	—†	$\varphi \triangledown \chi$	$\chi \vDash \varphi$	—†	$\chi \vDash \varphi$
$\varphi \bowtie \psi$	$\varphi \triangledown \chi$	$\varphi \vartriangle \chi$	$\varphi \bowtie \chi$	$\chi \vDash \varphi$	$\varphi \vDash \chi$	$\varphi \simeq \chi$

In cells marked with †, the fact that no relations hold in general can be seen by noting that, if ψ is a tautology, the given relations between it and φ and χ will hold no matter what sentences φ and χ are, so it is possible for φ and χ to be logically independent. And, in the cells marked with *, something similar holds in a case where ψ is absurd: the given relations between ψ and each of φ and χ will hold no matter what φ and χ are. There are various considerations which can be used to show that what is said in other cases is the most that can be said, but it is probably easiest just to confirm for yourself that no further truth values for φ and χ are ruled out by the given information about the relation of each to ψ.

Glen Helman 01 Aug 2013