1.4. General principles of deductive reasoning

1.4.0. Overview

All the deductive properties and relations of sentences can be seen as special cases of a single relation. We will look at this relation and also see how to study the full range of deductive logic by way of entailment and a couple of auxiliary ideas.

1.4.1. A closer look at entailment
Entailment will be at the heart of our study and we will begin by looking in some detail at a couple ways of formulating its definition.

1.4.2. Separation
It will be useful to have a special term for the kind of pattern of truth values that entailment rules out.

1.4.3. Content and coverage
The ideas of content and coverage can be extended to sets but in more than one way.

1.4.4. Relative exhaustiveness
Although entailment does not encompass all the concepts of deductive logic, there is a similarly defined relation that does.

1.4.5. A general framework
All the deductive properties and relations we will consider can be expressed in terms of relative exhaustiveness and expressed in a way that corresponds directly to definitions of them.

1.4.6. Reduction to entailment
Although relative exhaustiveness provides a way of thinking about deductive properties and relations, entailment is way that they are most naturally established, and we need to consider how this can be done.

1.4.7. Laws for entailment
The ideas behind the reflexivity and transitivity of implication provide the core of the general principles that hold for entailment by any number of premises.

1.4.8. Duality
The specific principles concerning $\top$ and $\bot$ display a kind of symmetry that we will also find in principles for other logical forms.

1.4.1. A closer look at entailment

In section 1.2 we looked at implication, the special case of entailment that applies to single-premised arguments or immediate inferences, and we looked at it in the context of other deductive relations between individual sentences. Now we will return to entailment in its full generality, as applying also to multiple-premised arguments, and consider also the full range of deductive properties and relations. These relations were each defined (in a way summarized in 1.2.8) as the rejection of the possibility of one or more patterns of truth values, and the properties of tautologousness and absurdity were each defined as the rejection of the possibility of a certain truth value for the sentence in question. Our task is now to extend this idea to collections of more than two sentences, and we will begin with entailment.

Entailment holds in the case of implication, $\varphi \vDash \psi$, when there is no possible world that separates φ from ψ. That definition can be stated more explicitly in either of the two following forms:

$\varphi \vDash \psi$	if and only if	there is no logically possible world in which ψ is false while φ is true
	if and only if	ψ is true in every logically possible world in which φ is true

These are not two different ideas, for the two statements to the right of if and only if say the same thing. Still, they provide different perspectives on implication. The second—which we will speak of as the *positive form* of the definition—is closely tied to the idea of a conditional guarantee of truth and thus to the reason why deductive inference is valuable. The first form—the *negative form*—makes the content of the guarantee especially clear—telling us that it is a guarantee against new error in moving from φ to ψ—and this form of definition will generally be the more useful when we try to prove things concerning the concept. The other deductive properties and relations we have discussed or will go on to discuss can be given analogous pairs of definitions, a negative form ruling out certain patterns of truth values and another form stating a more positive generalization.

The equivalence of the two forms of the definition reflects a feature of all generalizations. When a generalization is false, it is because of a *counterexample*. This is an example of the sort about which we generalize but that does not have the property we have said that all such things have, so a counterexample to the claim that all birds fly is a bird that does not fly. In the positive definition of implication, the generalization is about all possible worlds in which φ is true and such worlds are said to all have the property that ψ is true in them. A counterexample to such a generalization is then a world in which φ is true but ψ is not. The negative form of the definition then affirms the same generalization but by saying that no such counterexample exists. The added clarity of the negative definition reflects a rule of thumb applying to all generalizations: a good way to clarify a generalization is always to ask what sort of counterexample is being ruled out.

The analogous pair of definitions for entailment more generally characterize that relation as a guarantee against new error when adding the conclusion to a set of assumptions, or as a guarantee of the truth of the conclusion conditional on the truth of the assumptions:

$\Gamma \vDash \psi$	if and only if	there is no logically possible world in which ψ is false while all members of Γ are true
	if and only if	ψ is true in every logically possible world in which all members of Γ are true

These differ from the corresponding definitions of implication by replacing a reference to a single assumption φ by a reference to a set Γ. And since a set of sentences does not have a truth value, we need to speak of the truth of the assumptions by speaking of "all members of Γ" (a phrase whose significance we will return to in the next section).

Since we call an argument "valid" when its premises entail its conclusion, and validity is a good property for an argument to have, it is important to remember that validity is not all that we might ask of an argument. Compare a variant of the example in 1.1.3 with another argument having the same logical form:

All human beings are mortals	All dogs are reptiles
Socrates is a human being	Socrates is a dog
Socrates is a mortal	Socrates is a reptile

Since these arguments have the same form, they are equally valid and, indeed, valid for the same reason. But the second is clearly not a very good argument on other grounds. This is an instance of the general point that deductive logic is not concerned with the specific truth values of individual sentences (except in the special cases of tautologies or absurd sentences) but instead with ways in which the truth values of sentences are tied to one another. More specifically, the example emphasizes the fact that the relation of entailment rules out only one pattern of truth values, a false conclusion along with premises that are all true, and all other patterns can be found among deductive arguments. To take one further example, substituting god for human being in the first argument above shows that a valid argument may have a true conclusion even when its premises are all false.

1.4.2. Separation

The relation of implication can be characterized briefly by saying that $\varphi \vDash \psi$ when φ cannot be separated from ψ. We can do the same for entailment if we extend the idea of separation to say that a set Γ is *separated* from a sentence φ when the members of Γ are all true but φ is false. Then we can say that $\Gamma \vDash \varphi$ is true in just the cases where Γ cannot be separated from φ, and this way of thinking about entailment will reappear frequently throughout the course.

When we move from thinking of implication by a single premise to thinking of entailment by a set, we broaden our perspective to include inference from multiple premises. But reference to a set also admits the case that there are no premises or assumptions in question, for the set may be the *empty set*, which has no members. An argument that offers a conclusion without basing it on premises will be valid only if we have an unconditional guarantee that the conclusion is true—only if the conclusion is a tautology. We are provided no premises that could involve some falsehood, so any error in the conclusion is a new error, and a guarantee against new error must be a guarantee that it is true. And, if we look at validity in terms of content, we notice that there are no premises from which the content of the conclusion can be extracted, so it better have no content at all.

The same follows from our more formal definitions of entailment, but the idea may take a little while to get used to. First, let's adopt some notation for referring to the empty set. We will often use the sign "$\varnothing$", but we may also denote it by giving an empty list $\{\}$ of members. The latter notation fits with a way we will often write claims of entailment by the empty set: we might say that $\varnothing \vDash \varphi$, by saying $\vDash \varphi$—i.e., by writing an entailment sign with an empty list of assumptions to its left. Now the claim $\varnothing \vDash \varphi$ or $\vDash \varphi$ will hold just in case there is no possible world that separates $\varnothing$ from φ. But what is to separate the empty set from a sentence? Well, the possible world would need to be one in which φ was false and every member of $\varnothing$ was true. So we need to know what it takes for every member of $\varnothing$ to be true. The short answer is it takes nothing at all because $\varnothing$ has no members; that is, in every possible world it will be the case that every member of $\varnothing$ is true. For, to assert that every member of a set is true is to assert a generalization, so this assertion will be true unless there is a counterexample to the generalization. We can make that more explicit by restating the clause every member of Γ is true as no member of Γ is false, and the latter is clearly true when Γ has no members at all. We will run into generalizations about empty collections of things in other contexts, so it may help to have a label for them: they are often described as *vacuous gen-*

eralizations. The allusion to a vacuum is based on the idea that such generalizations are empty of content because there are no counterexamples available for them to rule out.

Since the condition for separation that concerns $\varnothing$ is vacuous, all that is necessary for $\varnothing$ to be separated from a sentence φ is for φ to be false. So, to say that $\varnothing$ cannot be separated from φ (i.e., that $\varnothing \vDash \varphi$ or $\vDash \varphi$) is to say that φ cannot be false, that it is a tautology. This gives us a simple notation for tautologousness: we can say that φ is a tautology by saying that $\vDash \varphi$. More importantly, it shows that we can study tautologousness by studying entailment because tautologousness is just a special case of entailment, namely, entailment by an empty set of assumptions.

The rest of this section will be devoted to showing how to study other deductive properties and relations by studying entailment, and a first step in doing that will be to extend the idea of separation further still. The properties and relations that I have been labeling "deductive" are ones that rule out certain patterns of truth values for the sentences they concern. An individual pattern of this sort will be a specification of truth values for certain sentences. It will make some (perhaps empty) set Γ of sentences all true and some set Σ all false, and we will say that in doing this it *separates* Γ from Σ. When Σ has a single member, this is just separation as we have been speaking of it, so we are now consideration that idea not only as a relation between a pair of sentences or between a set and a sentence but as one between any set and any other set. It can help, as it did in the case of entailment, to restate the requirement all members of Γ are true as no member of Γ is false and restate the requirement for Σ analogously. That is, separation in its full generality can be defined in either of the following ways, with the second one the clearest:

Γ is separated from Σ	if and only if	every member of Γ is true and every member of Σ is false
	if and only if	no member of Γ is false and no member of Σ is true

The requirement this places on a set is automatically satisfied when that set is the empty set $\varnothing$. That means that we can say:

Γ is separated from $\varnothing$ if and only if no member of Γ is false

$\varnothing$ is separated from Σ if and only if no member of Σ is is true

Either way, we can see that the empty set is bound to be separated from itself. That particular consequence is only a curiosity, but more interesting, and a taste of things to come, is the use of this broader notion of separation to describe the relations between pairs of sentences discussed in 1.2. In particular, $\varphi \vartriangle \psi$ if and only if $\{\varphi, \psi\}$, the set consisting of φ and ψ, cannot be separated from the empty set; for to say that $\{\varphi, \psi\}$ cannot be separated from the empty set is simply to say that φ and ψ cannot be made both true. Similarly, $\varphi \triangledown \psi$ if and only if the empty set cannot be separated from $\{\varphi, \psi\}$.

1.4.3. Content and coverage

In the case of implication, the idea of separation was tied to relations of content and coverage between sentences. When φ can be separated from ψ, we know that it is possible to have φ true and ψ false, so the content of φ does not include that of ψ and ψ does not cover every possibility that ψ does. On the other hand, when φ cannot be separated from ψ—i.e., when φ ⊨ ψ—we know that the content of ψ is part of the content of φ and that ψ covers any possibility that φ does.

These relations of content and coverage are closely tied to the ideas of extracting content or of having a conditional guaranteed of truth, so it is natural to extend them to entailment generally. But, to do this, we need to see how to apply the ideas of content and coverage to sets. We will actually apply them each in two different ways. Initially, we will use one approach for content and the other for coverage, but we will bring in the remaining two options shortly. The *cumulative content* of a set Γ—or, alternatively, its *content as a set of assumptions*—is set of all possible worlds in which every member of Γ is false. That means that the cumulative content of Γ includes the content of each of its members—it is the result of accumulating their individual contents—and this is the content we are interested in when asking what follows from a set of assumptions. The *coverage* of a set Γ *as a set of assumptions* consists of the remaining possible worlds, the ones that are not ruled out by any member of Γ and thus appear in the coverage of each of its members. We will describe this set also as the *shared coverage* of Γ since it consists of overlap in the coverages of the individual members of Γ. And when we use a deductively valid argument as support for our confidence in the coverage of a conclusion, what we know is that its coverage is at least as great as the shared coverage of the premises.

It's no surprise that the cumulative content of an empty set is empty, for there is nothing to accumulate. But it may seem surprising that its shared coverage is the full range of possibilities. This another instance of a vacuous generalization: the shared coverage of a set includes a possible world when every member of the set covers that world, and that is bound to be so when the set has no members. But it makes sense in its own right. We would expect that the shared coverage might decrease as we add members to a set since the new coverage in this sense will be limited to any overlap between what the set did cover and what is covered by a new member. But a set with one member φ has as its shared coverage just the coverage of φ, and it is the result of adding φ to the empty set. So the shared coverage of the empty set must have been at least that of φ, and that means it must have included every possible world because φ might be a tautology and have all possibilities in its coverage.

Of course it is possible to add up possibilities covered by the members of a set just as we have added up their contents. When we do this we will say that we are looking at the set as a set of *alternatives*. The idea is that, from this perspective, adding new members to the set can allow it to cover possibilities it did not before, so it adds new alternative ways of covering possibilities. So the *coverage* of a set Γ *as a set of alternatives* is the *cumulative coverage* of the set; it is the full range of possibilities covered by any of the set's members. When we look at a set *as a set of alternatives*, its *content* is the set of possibilities that are ruled out no matter what alternative we consider; that is, it is the set's *shared content*.

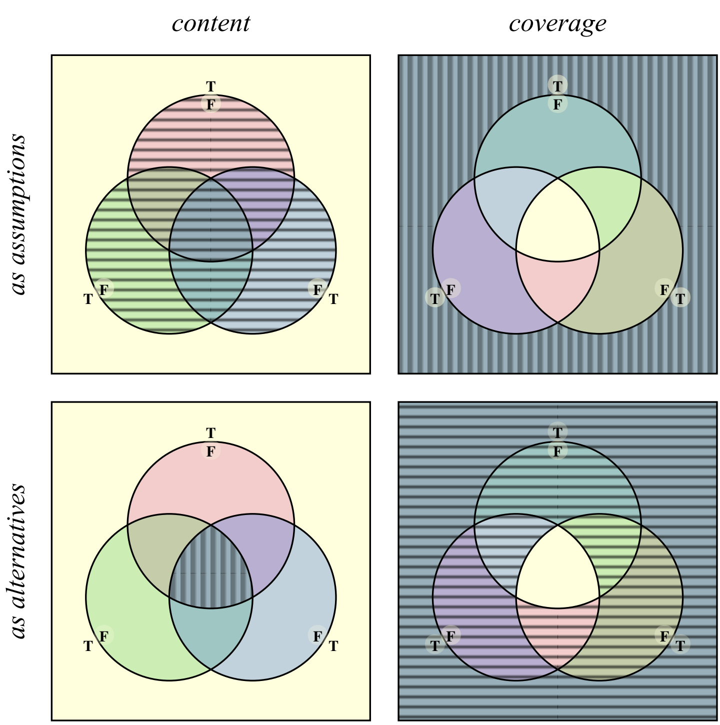

Fig. 1.4.3-1. A set of three sentences with the content of each sentence shaded on the left and its coverage shaded on the right. The top row shows, by hatching, the content and coverage of the set as a set of assumptions, and the bottom row shows its content and coverage as a set of alternatives. Cumulative content or coverage is hatched horizontally, and shared content or coverage is hatched vertically.

Figure 1.4.3-1 illustrates the application of ideas we been surveying to the case of a set of sentences with three members. Notice that the hatched areas side by side in each row are exactly opposite. The full range of possible worlds is divided between the content and coverage of any given sentence, and the same is true of the content and coverage of any set, provided we consider the set in both cases either as a set of assumptions or as a set of alternatives.

The application of these ideas to an empty set should now be no surprise. An empty set provides no alternatives, so its cumulative coverage is empty; and it has no members whose content must be respected in determining an overlap, so its shared coverage is the full range of possibilities.

Notice also that if we say that Γ is separated from Σ, we speak of a possibility that is in the shared coverage of Γ (because all of Γ's members are true) that is not in the cumulative coverage of Σ (because all of Σ's members are false). So if we reject the possibility of separation in this case, we say that the shared coverage of Γ is included in the cumulative coverage of Σ or, equivalently, that the cumulative content of Γ includes the shared content of Σ. In short, these two different perspectives on sets are built into the idea of separability that is at the root of all deductive properties and relations.

Glen Helman 14 Jul 2012

1.4.4. Relative exhaustiveness

We can use the idea of separation to generalize entailment to a relation between sets. And it is useful to do this because we can capture all the deductive properties and relations of sentences by using a relation that simply says that one set cannot be separated from another. That's what the following relation does:

$\Gamma \vDash \Sigma$	if and only if	there is no possible world in which all members of Γ are true and all members of Σ are false
	if and only if	in each possible world in which all members of Γ are true, at least one member of Σ is true

For reasons to be discussed shortly, when this relation holds, we will say that Γ *renders* Σ *exhaustive*, and we will use the phrase *relative exhaustiveness* to refer to this relation. We have reused the notation for entailment because entailment is the special case where the set Σ consists of a single sentence (just as implication is the special case of entailment where Γ consists of a single sentence).

While tautologousness is an unconditional guarantee of truth, entailment guarantees the truth of its conclusion only given the truth of a set of assumptions. Entailment is thus a guarantee of truth for a single sentence only given the conditions set out in the assumptions, and we can think about an analogous conditional guarantee that a set is exhaustive. Saying that Σ is exhaustive unconditionally tells us that its cumulative coverage includes all possibilities whatsoever. We can say that a set Σ is exhaustive *given* a set Γ when the cumulative coverage of Σ includes the shared coverage of Γ. For example, while the two alternatives The glass is full and The glass is empty are not jointly exhaustive, they are exhaustive given the assumption The glass is not partly full since it rules out all possibilities where they are both false. It is this sort of conditional exhaustiveness that is asserted by the relation above: relative exhaustiveness is exhaustiveness relative to a set of assumptions that limit the possibilities that the relatively exhaustive set must cover.

In cases of relative exhaustiveness that are not cases of entailment, what is rendered exhaustive is either a set of several alternatives or the empty set. In these cases, it does not make sense to speak of a conclusion, for when the set on the right has several members, these sentences need not be valid conclusions from the set that renders them exhaustive. Indeed, a jointly exhaustive pair of sentences will be rendered exhaustive by any set, but often neither member of the pair will be entailed by that set. This is particularly clear in the case of sentences like The glass is full and The glass is not full that are both jointly exhaustive and mutually exlcusive—i.e., that are contradictory. Although the set consisting of such pair is rendered exhaustive by any set, only an inconsistent set could entail both of these alternatives. So the term *conclusion* will be reserved for cases where there is a single sentence on the right of the sign $\vDash$. When there is more than one, we will speak of these sentences as *alternatives*. This term is also appropriate when the set on the right is empty, for then the claim being made is that no alternatives are needed to cover the possibilities where the assumptions are all true—that is, they form an inconsistent set.

A case $\Gamma \vDash \Sigma$ of relative exhaustiveness indicates relations of both content and coverage, but the ideas of content and coverage must be applied in different senses to Γ and Σ. The cumulative content of Γ is said to include the shared content of Σ. We wouldn't expect it to include more than this because the members of Σ may go off in different, even incompatible directions to cover the possibilities left open by the members of Γ. And the cumulative coverage of Σ is said to include all of the shared coverage of Γ, but only its shared coverage because the members of Γ may contribute in different ways to narrowing the range of possibilities that Σ is said to cover.

Glen Helman 11 Jul 2012

1.4.5. A general framework

It was noted in the last section that relative exhaustiveness does not merely generalize entailment and absolute exhaustiveness but encompasses all deductive properties and relations. It is not surprising that does so if these properties and relations are understood to all consist in guarantees that certain parterns of truth values appear in no possible world. For any claim there is no world where certain sentences Γ are true and other sentences Σ are false is a claim that $\Gamma \vDash \Sigma$. Of course, a given deductive property or relation may rule out a number of different patterns—i.e., rule out a number of different ways of distributing truth values among the sentences it applies to—but this just means that a deductive property or relation may consist of a number of different claims of relative exhaustiveness. In the case of the properties and relations we will consider, only equivalence and contradictoriness involve more than one claim of relative exhaustiveness.

The table below summarizes the deductive properties and relations that involve only one claim of relative exhaustiveness, and also shows the vocabulary we have used for various special cases. The first row covers cases where the number of assumptions is unspecified, with the next three concerning the cases of specific numbers of assumptions. Similarly, the first column places no constraints on the number of alternatives while the following three columns do. As a result, the first cell in the first row, the one for relative exhaustiveness encompasses all the rest. Notice, for example, that because tautologousness concerns a single alternative and no assumptions, it is a special case of both entailment and exhaustiveness.

alternatives

assumptions	*any no.*	*two*	*one*	*none*
any no.	$\Gamma \vDash \Sigma$	$\Gamma \vDash \psi, \psi'$	$\Gamma \vDash \psi$ *entails*	$\Gamma \vDash$ *inconsistent*
two	$\varphi, \varphi' \vDash \Sigma$	$\varphi, \varphi' \vDash \psi, \psi'$	$\varphi, \varphi' \vDash \psi$	$\varphi, \varphi' \vDash$ (or $\varphi \vartriangle \varphi'$) *mutually exclusive*
one	$\varphi \vDash \Sigma$	$\varphi \vDash \psi, \psi'$	$\varphi \vDash \psi$ *implies*	$\varphi \vDash$ *absurd*
none	$\vDash \Sigma$ *exhaustive*	$\vDash \psi, \psi'$ (or $\psi \triangledown \psi'$) *jointly exhaustive*	$\vDash \psi$ *tautologous*	$\vDash$

Since there are no alternatives in question, the ideas in the last column are really properties of sets of assumptions (just as those in the last row are properties of sets of alternatives). It does not make a claim about some sentence or set of sentences but about entailment itself, and the claim it makes is false. Since it concerns the case of no assumptions and no alternatives, it might be written more explicitly as the claim that $\varnothing \vDash \varnothing$. As was noted at the end of 1.4.2, the empty set $\varnothing$ is bound to be separated from itself, so $\varnothing \vDash \varnothing$ is bound to be false. And that's a result we should expect because this case of relative exhaustiveness offers an unconditional guarantee that some member of the empty set is true.

The ideas of separation and relative exhaustiveness also provide ways of extending to any set the idea of logical independence introduced in 1.2.8 in the case of a pair of sentences. First, let us look at this general idea of logical independence directly. We will say that a set Γ of sentences is *logically independent* when every way of assigning a truth value to each member of Γ is exhibited in at least one possible world. This is the same as saying that, for every part of the set (counting both the empty set and the whole set Γ as parts of Γ), it is possible to separate that part from the rest of the set. When the set has two members, this is the same as the earlier idea of logical independence; and when the set $\{\varphi\}$ containing a single sentence φ is logically independent in this sense, it is neither a tautology or absurd—i.e., it is logically contingent in the sense defined in 1.2.5.

Relative exhaustiveness provides an alternative way of describing logical independence of this general sort. For, when the sentences in a set are *not* independent, not every way of dividing them into a set of true sentences and a set of false sentences is logically possible. And when some way of dividing them is not possible, the set contains at least one pair of non-overlapping subsets Γ and Σ such that $\Gamma \vDash \Sigma$. And, of course, if the set contains such a pair, its members are not logically independent. So the members of a set are logically independent when the relation of relative exhaustiveness never holds between non-overlapping subsets. (It always holds between sets that overlap because then there clearly is no way of separating one set from the other.)

When a set is logically independent, each member is contingent and any two of its members are logically independent, but the contingency of members and the independence of pairs does not by itself imply that the set as a whole is logically independent. For example, assume that the sentences X is fast, X is strong, X has skill, and X has stamina form an independent set. Then the sentences

X is fast and strong	X has skill and stamina	X is fast and has stamina

are each contingent, and any two of them can be seen to be independent. However, the first two taken together entail the third, so these three more complex sentences do not form an independent set. This also shows that, while it is natural to speak of the members of the set as independent, independence in this sense is really a property of the set as a whole. For we can say that the two sentences X is fast and strong and X is fast and has stamina are independent as a pair, but adding the third gives us a set whose members do not count as independent with respect to the other two.

Glen Helman 11 Jul 2012

1.4.6. Reduction to entailment

Relative exhaustiveness generalizes entailment by allowing cases in which we have, instead of a single conclusion, multiple alternatives or none at all. To express the ideas captured by relative exhaustiveness in terms of entailment, we need to add ways of capturing each of these added cases.

When a claim of relative exhaustiveness offers no alternatives, it asserts the inconsistency of the assumptions; and that comes to the same thing as entailing the specific absurdity ⊥. That is, we can state the following:

INCONSISTENCY *VIA* ABSURDITY. $\Gamma \vDash$ (i.e., $\Gamma \vDash \varnothing$) if and only if $\Gamma \vDash \bot$.

This law holds because rendering exhaustive the empty set and entailing ⊥ both offer conditional guarantees of a truth that cannot exist, so each has the effect of ruling out the possibility of meeting the conditions of the guarantee. Alternatively (but equivalently) we can note that ⊥, since it cannot be true, covers the same possibilities as the empty set, so, like the empty set, it covers the possibilities in the shared coverage of Γ only when that is empty. And to say that the shared coverage of a set is empty is to say that its members cannot all be true.

To express the idea of rendering exhaustive multiple alternatives using entailment we need help from the concept of contradictoriness. Contradictoriness comes in here because having an exception in a guarantee—which is what an added alternative provides—comes to the same thing as having the contradictory of this exception as a condition. For example, the guarantee The product will funciton for three years unless it is abused is equivalent to The product will function for three years if it hasn't been abused, and the guarantee The product will function for three years if it is serviced regularly is equivalent to The product will function for three years unless it is not serviced regularly. In the first case we move from an exception to a condition, and in the second we move in the opposite direction. To make this intuitive point more formally, note first that when sentences are contradictory, they always have opposite truth values. So making one true comes to the same thing as making the other false, and that means that contradictory sentences play opposite roles when one set is being separated from another. More specifically, if φ and $\varphi^{\bowtie}$ are contradictory sentences, then

$$\Gamma \text{ is separated from } (\Sigma \text{ together with } \varphi)$$
$$\textit{if and only if}$$
$$(\Gamma \text{ together with } \varphi^{\bowtie}) \text{ is separated from } \Sigma$$

because each of these separations requires that φ be made false and $\varphi^{\bowtie}$ be made true. Since a claim of relative exhaustiveness asserts that a separation is not possible, having a sentence as an alternative comes to the same thing as having a sentence contradictory to it as an assumption; that is,

$$\text{if } \varphi \text{ and } \varphi^{\bowtie} \text{ are contradictory, then } \Gamma \vDash \varphi, \Sigma \text{ if and only if } \Gamma, \varphi^{\bowtie} \vDash \Sigma$$

If we apply this idea repeatedly, we can move a set of alternatives to the left of the turnstile, and the direct justification for doing that is the same: having a collection of the sentences in the right comes to the same thing, as far as separation is concerned, as having on the left sentences contradictory to the members of the collection. If we assume there is no limit on the number of times this can be done, we get the following law:

ALTERNATIVES *VIA* CONTRADICTORY ASSUMPTIONS. Let $\Delta^{\bowtie}$ be the result of replacing each member of Δ by a sentence contradictory to it. Then $\Gamma \vDash \Delta, \Sigma$ if and only if $\Gamma, \Delta^{\bowtie} \vDash \Sigma$.

In short, we can remove alternatives if we put sentences contradictory to them among the assumptions.

The laws we have seen give us two approaches to restating claims of relative exhaustiveness as entailments. A claim with no alternatives—i.e., a claim of inconsistency—can be turned into an entailment by adding $\bot$ as the conclusion. And we may replace any alternatives by assumptions contradictory to them to reduce multiple alternatives to a single conclusion. The two may be combined by replacing all alternatives by contradictory assumptions and then adding $\bot$ as conclusion.

The following table uses these two approaches to restate all the deductive properties shown in the table of the last subsection:

assumptions \ *alternatives*	*any no.*	*two*	*one*	*none*
any no.	$\Gamma \vDash \Sigma$ $\Gamma, \Sigma^{\bowtie} \vDash \bot$	$\Gamma \vDash \psi, \psi'$ $\Gamma, \psi^{\bowtie} \vDash \psi'$	*entails* $\Gamma \vDash \psi$ (*same*)	*inconsistent* $\Gamma \vDash$ $\Gamma \vDash \bot$
two	$\varphi, \varphi' \vDash \Sigma$ $\varphi, \varphi', \Sigma^{\bowtie} \vDash \bot$	$\varphi, \varphi' \vDash \psi, \psi'$ $\varphi, \varphi', \psi^{\bowtie} \vDash \psi'$	$\varphi, \varphi' \vDash \psi$ (*same*)	*mutually excl.* $\varphi, \varphi' \vDash$ $\varphi, \varphi' \vDash \bot$
one	$\varphi \vDash \Sigma$ $\varphi, \Sigma^{\bowtie} \vDash \bot$	$\varphi \vDash \psi, \psi'$ $\varphi, \psi^{\bowtie} \vDash \psi'$	*implies* $\varphi \vDash \psi$ (*same*)	*absurd* $\varphi \vDash$ $\varphi \vDash \bot$
none	*exhaustive* $\vDash \Sigma$ $\Sigma^{\bowtie} \vDash \bot$	*jointly exh.* $\vDash \psi, \psi'$ $\psi^{\bowtie} \vDash \psi'$	*tautologous* $\vDash \psi$ (*same*)	$\vDash$ $\vDash \bot$

The natural statement of the property or relation in terms of relative exhaustiveness is shown first, followed by a statement in terms of entailment if that is different. The alterations are made in the same way for each column. In the last column, $\bot$ is added to get a conclusion; in the second, the alternative ψ is removed and its contradictory $\psi^{\bowtie}$ is added as an assumption; and, in the first, the set of alternatives Σ is replaced by $\bot$ and the contradictories $\Sigma^{\bowtie}$ of the members of Σ are added as assumptions.

There are other says of stating most of these ideas in terms of entailment, absurdity, and contradictoriness. Any time $\bot$ appears as the conclusion and there is at least one assumption, $\bot$ could be replaced as the conclusion by a sentence contradictory to some assumption, which is then dropped from the assumptions. That is, $\Gamma, \varphi \vDash \bot$ if and only if $\Gamma \vDash \varphi^{\bowtie}$. And whenever $\bot$ is not the conclusion, it could be made the conclusion if the a sentence contradictory to the previous conclusion is added to the assumptions—i.e., $\Gamma \vDash \varphi$ if and only if $\Gamma, \varphi^{\bowtie} \vDash \bot$. So, in particular, saying that $\varphi^{\bowtie} \vDash \psi$ comes to the same thing as saying that $\varphi^{\bowtie}, \psi^{\bowtie} \vDash \bot$, which comes to the same thing as saying that $\psi^{\bowtie} \vDash \varphi$. In particular, the claim of relative exhaustiveness beginning the following list can be restated as any of the claims of entailment after it:

the temperature is extreme ⊨ it's very hot, it's very cold
the temperature is extreme, it's not very hot ⊨ it's very cold
the temperature is extreme, it's not very cold ⊨ it's very hot
the temperature is extreme, it's not very hot, it's not very cold ⊨ ⊥

(This is an instance of the third row, second column of the table.)

It may seem pointless to define the relation of contradictoriness in terms of entailment, as is done in the last row of the table, since we need to use the idea of contradictoriness in order to do this. But the definition does mean that, once we know a single sentence contradictory to a given sentence, we can say what other sentences are contradictory to it using only the ideas of entailment and absurdity.

1.4.7. Laws for entailment

Most of the laws of deductive reasoning we will study will be generalizations about specific logical forms that will be introduced chapter by chapter, but some very general laws can be stated at this point. We have already seen some of these. We have just seen the laws tying inconsistency to Absurdity alternatives to assumptions. And the principles of reflexivity and transitivity for implication discussed in 1.2.3 can be generalized to provide basic laws for entailment and relative exhaustiveness. However, we will look only at the case of entailment.

Two basic laws suffice to capture the basic properties of entailment considered in its own right:

LAW FOR PREMISES. *Any set of assumptions entails each of its members.* That is, $\Gamma, \varphi \vDash \varphi$ (for any sentence φ and any set Γ).

CHAIN LAW. *A set of assumptions entails anything that is entailed by things the set entails.* That is, if $\Gamma \vDash \varphi$ for each assumption φ in Δ and $\Delta \vDash \psi$, then $\Gamma \vDash \psi$ (for any sentence ψ and any sets Γ and Δ).

Think about the relation which holds between sets Γ and Δ when Γ entails all members of Δ. Although a relation between sets, this is different from relative exhaustiveness because it says that the cumulative content of Δ (and not merely its shared content) is included in the cumulative content of Γ. That is, we are looking at both Γ and Δ as sets of assumptions. The laws above tell us that this relation is both reflexive and transitive. For the law for premises tells us that any set entails every member of itself. And, if Γ entails every member of Δ and Δ entails every member of Ξ, then Γ also entails every member of Ξ by the chain law. (On the other hand, relative exhaustiveness is neither reflexive nor transitive.)

The two principles above have as a consequence two further principles that concern the addition and subtraction of assumptions and will play an important role in our study of entailment:

MONOTONICITY. *Adding assumptions never undermines entailment.* That is, if $\Gamma \vDash \varphi$, then $\Gamma, \Delta \vDash \varphi$ (for any sets Γ and Δ and any sentence φ).

LAW FOR LEMMAS. *Any assumption that is entailed by other assumptions may be dropped without undermining entailment.* That is, if $\Gamma, \varphi \vDash \psi$ and $\Gamma \vDash \varphi$, then $\Gamma \vDash \psi$ (for any sentence φ and set Γ).

Each of these principles is based on both the law for premises and the chain law. In the case of the first, the law for premises tells us that Γ together with Δ entails every member of Γ alone, so if we also know that $\Gamma \vDash \varphi$, the chain law tells us that $\Gamma, \Delta \vDash \varphi$. The assumption of the second principle that $\Gamma \vDash \varphi$ combines with the law for premises to tell us that Γ entails every member of the set of assumptions consisting of Γ together with φ, and the chain law then tells us that Γ alone entails anything ψ that is entailed by this enlarged set of assumptions.

The term lemma can be used for something that we conclude not because it is of interest in its own right but because it helps us to draw further conclusions. The second law tells us that if we add to our premises Γ a lemma φ that we can conclude from them, anything ψ we can conclude using the enlarged set of premises can be concluded from the original set Γ.

The idea behind the law of monotonicity is that adding assumptions can only make it harder to find a possible world that separates the assumptions from the conclusion, so, if no possible world will separate Γ from φ, we can be sure that no world will separate from φ the larger set of assumptions we get by adding some further assumptions Δ. The term monotonic is applied to trends that never change direction. More specifically, it is applied to a quantity that does not both increase and decrease in response to changes in another quantity. In this case, it reflects the fact that adding assumptions will never lead to a decrease in the range of conclusions that are valid.

It is a distinguishing characteristic of deductive reasoning that a principle of monotonicity holds. For, when reasoning is not risk free, additional data can show that a initially well-supported conclusion is false and do so without undermining the original premises on which the conclusion was based. But then, if such further data were added to the original premises, the resulting enlarged set of assumptions would no longer support the conclusion. This means that risky inference is, in general, *non-monotonic* in the sense that additions to the premises can reduce the set of conclusions that are justified.

This is true of inductive generalization and of inference to the best explanation of available data, but the term non-monotonic is most often applied to another sort of non-deductive inference, an inference in which features of typical or normal cases are applied when there is no evidence to the contrary. One standard example is the argument from the premise Tweety is a bird to the conclusion Tweety flies. This conclusion is reasonable when the premise exhausts our knowledge of Tweety; but the inference is not free of risk, and the conclusion would no longer be reasonable if we were to add the premise Tweety is a penguin.

The law for premises and the chain law can be shown to give a complete account of the general laws of entailment in the sense that any relation between sets of sentences and sentences that obeys them is an entailment relation for some set of possibile worlds and assignment of truth values to sentences in each world. But this is not to say that they provide a complete general account of deductive properties and relations, because our definitions of many of these in terms of entailment also used the ideas of contradiction and absurdity. The laws providing for inconsistency *via* absurdity and for alternatives *via* assumptions govern these ideas but they were stated for relative exhaustiveness rather than entailment. In the next subsection, we will look at laws for $\bot$. Laws for contradiction will be considered by way of the account of negation in 3.2.1. The basic idea is that a pair of contradictory sentences each exclude the other and are the weakest way of doing that in the sense that each is entailed by any set of assumptions that excludes the other.

1.4.8. Duality

In the context of relative exhaustiveness all that need be said about the logical properties of Tautology $\top$ and Absurdity $\bot$ is that Tautology is a tautology (i.e., $\vDash \top$) and that Absurdity is absurd (i.e., $\bot \vDash$). The first of these makes sense for entailment and, together with the basic laws of entailment, provides the basis for the sort of laws for $\bot$ we will consider shortly. However, it is the latter laws that we will focus on since they state the role of $\top$ in entailment. And, in the case of $\bot$, saying merely that it is absurd tells us nothing from the point of view of entailment since that is to say only that $\bot \vDash \bot$.

Tautology $\top$ is entailed by any set of premises (the empty set included) because it cannot go beyond the information contained in any set of sentences; and, for the same reason, the presence of $\top$ among the premises of an argument contributes nothing to the argument's validity. These two ideas can be expressed more formally in the following laws.

LAW FOR $\top$ AS A CONCLUSION. $\Gamma \vDash \top$ (for any set Γ).

LAW FOR $\top$ AS A PREMISE. $\Gamma, \top \vDash \varphi$ if and only if $\Gamma \vDash \varphi$ (for any set Γ and sentence φ).

Although they are stated for $\top$, these laws will hold for all tautologies since they hold simply in virtue of the proposition expressed by $\top$.

These laws are different in character from the ones consider in the last subsection because they concern the logical properties of a specific sort of sentence rather than the general principles governing logical relations. They are also a first sample of a common pattern in the laws of deductive reasoning that we will consider. Entailment is so central to deductive reasoning that an account of the role of a kind of sentence in entailment as a conclusion and as a premise will usually tell us all we need to know about it.

A simple law describes the role of absurdities as premises. We state it for the specific absurdity $\bot$.

LAW FOR $\bot$ AS A PREMISE. $\Gamma, \bot \vDash \varphi$ (for any set Γ and sentence φ).

An argument with an absurdity among its premises is valid by default. Since its premises cannot all be true, there is no risk of *new* error no matter what the conclusion is. There is no law restating the significance of having $\bot$ as a conclusion because that is simplest way we have of using entailment to say that a set of assumptions is inconsistent.

Although entailment will be our focus, it is enlightening to consider analogues for relative exhaustiveness of the laws just stated. In particular, we can state a law for $\bot$ as an alternative in the context of relative exhaustiveness, and all the properties of $\top$ and $\bot$ take a particularly symmetric form when stated in terms of that relation.

	as a premise	*as an alternative*
Tautology	if $\Gamma, \top \vDash \Sigma$, then $\Gamma \vDash \Sigma$	$\Gamma \vDash \top, \Sigma$
Absurdity	$\Gamma, \bot \vDash \Sigma$	if $\Gamma \vDash \bot, \Sigma$, then $\Gamma \vDash \Sigma$

That is, while $\top$ contributes nothing as a premise and may be dropped, it is enough for a claim of relative exhaustiveness to hold that it be an alternative (no matter how small the set Γ of premises or the set Σ of other alternatives). And while it is enough to have $\bot$ as a premise (no matter how small the set of alternatives is), it contributes nothing as an alternative and may be dropped.

Notice that the converses of the principles at the upper left and lower right hold by monotonicity because they are just the addition of a premise in one case and an alternative in the other. If we take the if and only if principle that results from adding the converse to the lower right and consider a case where Σ is empty, we get

$$\Gamma \vDash \bot \text{ if and only if } \Gamma \vDash$$

This is the principle for relative exhaustiveness that lies behind the law providing inconsistency *via* Absurdity of 1.4.6. The moral is that our use of $\bot$ as a conclusion to define inconsistency in terms of entailment really involves the same idea as the principle for $\bot$ as an alternative that may be stated for relative exhaustiveness.

The symmetry exhibited by the set of principles in the table above might be traced to the fact that $\top$ and $\bot$ are contradictory since then having one as an assumption comes to the same thing as having the other as an alternative according to the law of 1.4.6 providing alternatives *via* assumptions. However, there is a more general idea behind this symmetry that will apply also to cases where sentences are not contradictory.

The essential difference between the lower left and upper right in the table above lies in interchanging $\bot$ and $\top$ and, at the same time, interchanging premises and alternatives. And the same is true of the upper left and lower right. That is, if we apply this transition to the lower left, we get

$$\Sigma \vDash \Gamma, \top$$

and that differs from the upper right only in the order of the alternatives and the exchange of Σ for Γ. And neither of these differences is important. Alternatives function only as a set, so the order in which they are listed does not matter. And, since each of Γ and Σ could be any set, exchanging these labels does not alter the content of the principle. Either way, we say that it is enough to have ⊤ as an alternative no matter what premises and what further alternatives we have. The possibility of the sort of transformation used to get from the lower left to the upper right can be expressed by saying that ⊤ and ⊥ on the one hand and premise (or assumption) and alternative on the other constitute pairs of *dual* terms. We will run into other pairs of terms later that combine with these pairs in an even broader sort of duality.

Glen Helman 11 Jul 2012

1.4.s. Summary

1 Entailment may be defined in two equivalent ways, negatively as the relation that holds when the conclusion is false in no possible world in which all the premises are true or positively as the relation which holds when the conclusion is true in all such worlds. The negative form has the advantage of focusing attention on the sort of possible world that serves as a counterexample to a claim of entailment. The positive form characterizes a relation of entailment as a conditional guarantee of the truth of the conclusion, a guarantee conditional on the truth of the premises.

2 The requirements for a world to serve as a counterexample to entailment suggest the general idea of separating one set from another by making all members of the first true and all members of the second false.

3 When we extend the ideas of content and coverage to sets, we can do this by adding up the content or coverage of individual members—yielding cumulative content or cumulative coverage—or by selecting out the part of the content or coverage of that all members share—yielding shared content or shared coverage. When we look at a set as a set of assumptions, we think of added assumptions as narrowing down possibilities, so we are interested in content in the cumulative sense and coverage in the shared sense. On the other hand, when we look at a set as a set of alternative ways of covering possibilities, it is coverage for which the cumulative sense is appropriate, and we are interested in content in the shared sense.

4 The idea of separation enables us to define a relation of relative exhaustiveness between sets: one set renders another exhaustive when there is no possible world that separates the the first from the second. We will extend the notation for entailment to express this relation between sets Γ and Σ as $\Gamma \vDash \Sigma$. Entailment is the special case of this where Σ has only one member. When Σ has more than one member, its members count as alternatives because a relation of relative exhaustiveness provides a conditional guarantee only that at least one member of the second set it true.

5 Since a set of alternatives can have more than one member or be empty, relative exhaustiveness encompasses all the deductive properties and relations we have considered (as well as an extension of the idea of joint exhaustiveness to any set of sentences). The way a property or relation is expressed using relative exhaustiveness is tied directly to the negative form of the definition of the property or relation. When no relation of relative exhaustiveness holds no matter how a set is divided into two parts, all patterns of truth val-

ues for its members are possible and the set is logically independent.

6 Definitions in terms of relative exhaustiveness can be converted into definitions in terms of entailment by replacing empty sets of alternatives with ⊥ and reducing the size of multiple sets of alternatives by replacing members by adding assumptions that are contradictory to them (using the law for alternatives *via* contradictory assumptions).

7 Entailment obeys analogues to the principles of reflexivity and transitivity for implication. In the case of reflexivity, the analogy is with the law for premises; and, in the case of transitivity, it is with the chain law. Taken together, these principles yield all laws of entailment. Two principles for entailment that follow from them—monotonicity and the law for lemmas—state conditions under which we may add and drop assumptions. The second principle licenses the use of lemmas, valid conclusions that are of interest only as premises in further arguments. The first tells us that entailment is monotonic in the sense that it will never stop holding because of additions to the set of assumptions. This principle is significant in distinguishing entailment from other forms of good inference, whose riskiness means that they are non-monotonic (because adding information telling us that the risk does not pay off will undermine their quality).

8 The laws describing the behavior of ⊤ and ⊥ in the context of relative exhaustiveness exhibit a kind of symmetry that we will see in other laws later. The sentences ⊤ and ⊥ are dual as are the terms premise and alternative (or the left and right of an turnstile) in the sense that replacing each such term in a law by the one dual to it will produce another law.

Glen Helman 11 Jul 2012

1.4.x. Exercise questions

1. Any claim that a deductive relation holds can be stated as one or more claims that one set of sentences cannot be separated from another. (i) Restate each of the following claims in that way, and (ii) explicitly describe the sort of possibility that would separate the first of the sets from the second and is thus ruled out by claiming that the deductive relation holds. Nonsense words have been used to help you think to think how a possibility would be described without worrying whether that possibility could really occur.

For example, the claim that the sentences The widget plonked and The widget plinked are equivalent can be restated by saying that (i) the set consisting of the first sentence cannot be separated from the set consisting of the second sentence and vice versa. That is, (ii) it rules out any possibility in which the widget plonked but did not plink and any possibility in which the widget plinked but did not plonk.

a. The gizmo is a widget and The gizmo is a gadget are mutually exclusive

b. The gizmo is a widget and The gizmo is a gadget are jointly exhaustive

c. The widget plinked is a tautology

d. The widget plonked is absurd

e. The widget was a gadget renders exhaustive the alternatives The widget plinked and The widget plonked

f. The widget was a gizmo, The widget plinked, and The widget plonked are inconsistent

2. The basic law for relative exhaustiveness can be used not only to replace alternatives by assumptions but also to replace assumptions by alternatives. For example, the claim that The widget is blue entails The widget is colored can be restated to say (i) The widget is blue and The widget is not colored are inconsistent, (ii) The widget is not blue and The widget is colored form an exhuastive set, or (iii) The widget is not colored entails The widget is not blue.

In the following, you will be asked to restate some statements of deductive relations by replacing alternatives with assumptions or assumptions with alternatives. You may add or remove ordinary negation to state the contradictories of sentences.

a. Restate the following as a claim of entailment: The gadget is red and The gadget is green are mutually exclusive

b. Restate the following as a claim of entailment: Someone is in the auditorium and There are empty seats in the auditorium are jointly exhaustive

c. Restate the following as a claim of absurdity: A widget is a widget is a tautology

d. Restate the following as a claim of tautologousness: A widget is a gadget is absurd

e. Restate the following as a claim of inconsistency: The widget is a gadget or gizmo and The widget is not a gadget entail The widget is a gizmo

f. Restate the following so that each assumption is replaced by an alternative and each alternative by an assumption: The widget has advanced and The widget has plonked render exhaustive the alternatives The widget has finished the task and The widget has broken

Glen Helman 11 Jul 2012

1.4.xa. Exercise answers

1. **a.** (i) The set consisting of The gizmo is a widget and The gizmo is a gadget cannot be separated from the empty set; that is, (ii) there is no possibility of the gizmo being both a widget and a gadget.

b. (i) The empty set cannot be separated from the set consisting of The gizmo is a widget and The gizmo is a gadget; that is, (ii) there is no possibility of the gizmo being neither a widget nor a gadget

c. (i) The empty set cannot be separated from the set consisting of only The widget plinked; that is, (ii) there is no possibility that the widget did not plink

d. (i) The set consisting of only The widget plonked cannot be separated from the empty set; that is, (ii) there is no possibility that the widget plonked

e. (i) The set consisting of only The widget was a gadget cannot be separated from the set consisting of The widget plinked and The widget plonked; that is, (ii) there is no possbility that the widget was a gadget while not either plinking or plonking.

f. (i) The set consisting of The widget was a gizmo, The widget plinked, and The widget plonked cannot be separated from the empty set; that is, (ii) there is no possibility that the widget was a gizmo and both plinked and plonked

2. **a.** The gadget is red entails The gadget is not green (*or:* The gadget is green entails The gadget is not red)

b. The auditorium is empty entails There are empty seats in the auditorium (*or:* There are no empty seats in the auditorium entails The auditorium is not empty)

c. A widget is a not widget is absurd

d. A widget is a not gadget is a tautology

e. The widget is a gadget or gizmo, The widget is not a gadget, and The widget is not a gizmo are inconsistent

f. The widget has not finished the task and The widget has not broken render exhaustive The widget has not advanced and The widget has not plonked