2. Conjunctions

2.1. *And*: adding content

2.1.0. Overview

In this chapter, we will study the logical properties of the English word and and certain related expressions. Along the way we will encounter some general ways of approaching the study of logical properties that will serve us in later chapters, too. For the next several chapters, will be interested in sentences that are formed from other sentences; the operations used to do this are known as *connectives*.

2.1.1. A connective
We begin with *conjunction*, a connective that enables us to combine the content of a pair of sentences.

2.1.2. A truth function
The meaning of conjunction can be given by specifying the truth value of a conjunction in terms of the truth values of the sentences that were combined.

2.1.3. Conjunction in English
Although conjunction is most closely associated with the word and, there are a number of ways of expressing it in English.

2.1.4. Limits on analysis
On the other hand, even the appearance of and is not a sure sign that a sentence may be analyzed as a conjunction.

2.1.5. Multiple conjunction
The operation of forming a sentence from sentences can be repeated. We will look at this sort of iteration in the case of conjunction.

2.1.6. Some sample analyses
We will then apply these ideas to analyze several examples.

2.1.7. Logical forms
And we will look in more general terms at the relation of logical forms to actual sentences.

2.1.8. Interpretations
Finally, we will introduce some ways of talking about the relation between abstract forms and the meanings of sentences.

2.1.1. A connective

We are interested in logical forms as a way of stating general laws of entailment. Let us begin by looking at cases of entailment that seem to involve the word and. Here is an example:

That bear is large and edgy $\vDash$ That bear is large

In attempting to understand any fact, it is useful to collect related facts. One way to search for related facts about entailment is to look for cases involving sentences similar in grammatical form to those above. If we follow this route, we run into entailments like this:

That car is cheap and reliable $\vDash$ That car is cheap

And we will eventually hit upon a general pattern like this:

a is P and Q $\vDash$ a is P.

Although we will move on to more general patterns, any pattern that abstracts from particular words makes the label "formal logic" appropriate.

If we look a little farther afield, we also find examples like

It was hot and there was a storm before dark $\vDash$ It was hot,

which follows the pattern

φ and $\psi \vDash \varphi$.

This pattern can be seen to operate also in examples of the first group if we paraphrase them, transforming

That bear is large and edgy,

for instance, into

That bear is large and that bear is edgy.

In applying a pattern by first paraphrasing, we treat a sentence as having a form that is hidden by its surface appearance. Much of our analysis of logical form will involve this sort of transition.

Assuming we are willing to apply the second pattern by way of a paraphrase, we have a fairly general law of inference in which the word and plays a key role. If we look at what this role involves, we see that and marks a particular sort of *compound* sentence formed of *component* sentences, one that we will label a *conjunction*. So the word and is a sign for an operation that forms conjunctions. We will call an operation that forms compound sentences out of component sentences a *connective*, and we will refer to the connective we are considering here as *conjunction*, marking it with the sign $\wedge$ (one of whose names is *logical and*). (The use of the term conjunction for both the operation of conjoining and the compound that results from it may seem confusing, but it follows a pattern that is used fairly often in English—as when the word *distribution* is used both for the act of distributing and for its result.) It will often be convenient to employ a further related term and refer to the components of a conjunction as its *conjuncts*.

Using these ideas, we can express our *analysis* of That bear is large and edgy as

$$\text{That bear is large} \wedge \text{that bear is edgy},$$

and we can express our principle of entailment as

$$\varphi \wedge \psi \vDash \varphi.$$

This symbolic notation can save space, but it is often convenient to use English to mark conjunction. When we do this, we will use the construction both … and … and write it (as done here) using a special type. So the principle above could be stated as

$$\text{both } \varphi \text{ and } \psi \vDash \varphi.$$

(The reason for using the particle both in addition to and will be discussed later.)

At this point, we have reached a stage like that reached by a physicist who recognizes pressure, temperature, and volume as physical quantities and has formulated a law relating them but who does not know why the law holds. That is, we have a generalization about entailment that we can apply in special cases, but we cannot say why this generalization is true. What is it about conjunction that makes this sort of entailment work?

We can find an answer by again scaring up some more facts. Notice, for example, that the entailment that got us started is matched by a second.

$$\text{That bear is large and edgy} \vDash \text{That bear is edgy}.$$

Moreover, we can see not only that the sentence That bear is large and edgy entails each of the two sentences That bear is large and That bear is edgy but also that it is entailed in turn by the two taken together.

If we abbreviate the longer sentence by B and the two shorter sentences as L and N, respectively, we have collected the following facts:

$$B \vDash L$$
$$B \vDash N$$
$$L, N \vDash B$$

And checking other cases of conjunction would show us that these are instances of three general laws.

$$\varphi \wedge \psi \vDash \varphi$$
$$\varphi \wedge \psi \vDash \psi$$
$$\varphi, \psi \vDash \varphi \wedge \psi.$$

Or, using English to express the forms,

$$\text{both } \varphi \text{ and } \psi \vDash \varphi$$
$$\text{both } \varphi \text{ and } \psi \vDash \psi$$
$$\varphi, \psi \vDash \text{both } \varphi \text{ and } \psi.$$

So far, all we have done is to accumulate more general laws, but it is often easier to understand a larger number of facts because a pattern can begin to emerge.

2.1.2. A truth function

We can begin to provide a single account of the laws outined in 2.1.1 by recalling our definition of entailment. In positive form, it says that a set Γ entails a sentence φ if and only if φ is **T** in every possible world in which each member of Γ is **T**. Restating our three laws in these terms, we have

φ is **T** in every possible world in which φ ∧ ψ is **T**
ψ is **T** in every possible world in which φ ∧ ψ is **T**
φ ∧ ψ is **T** in every possible world in which both φ and ψ are **T**

In short, φ ∧ ψ is true in a possible world if and only if both φ and ψ are true.

This means that the truth value of the compound φ ∧ ψ is determined by the truth values of the components φ and ψ, a fact we can express in the *truth table* below.

φ	ψ	φ ∧ ψ
T	T	T
T	F	F
F	T	F
F	F	F

This table shows the contribution of conjunction to the truth conditions of compound sentences formed using it, for it tells us how to determine the truth value of a conjunction φ ∧ ψ in any possible world once we know the truth values of the conjuncts φ and ψ. And the table also shows what lies behind the general laws of entailment that led us to it: it is because conjunction makes this sort of contribution to the meaning of sentences that those laws hold. A particular way of associating an output truth value with input truth values is a *truth function*, so we can say that conjunction expresses a truth function, and we can say the laws of entailment stated above reflect the character of the truth function that conjunction expresses.

It is worth pausing a moment to look at the way in which the proposition that is expressed by a conjunction is related to the propositions that are expressed by its components. Since a conjunction is false whenever either component is false, it rules out any possibility ruled out by either component; and, since the possibilities ruled out are an indication of the information a sentence contains, we can say that a conjunction contains all the information contained in its components. This means that the effect of conjunction is to add up informational content. Now, more information means fewer possibilities left open and, looking at the table in these terms, we see that a conjunction leaves open only the possibilities left open by both components. The range of possibilities it leaves open is the region in the full space of possibilities where the ranges of possibilities left open by the two components overlap.

For example, the sentence The number shown by the die is odd and less than 4 can be analyzed as the conjunction The number shown by the die is odd ∧ the number shown by the die is less than 4. The first component rules out possibilities where the die shows 2, 4, or 6 and the second rules out possibilities where it shows 4, 5, or 6. The conjunction rules out all these possibilities—that is, any possibility where the die shows 2, 4, 5, or 6. Looking at things in terms of the possibilities left open, the first component leaves open those where the die shows 1, 3, or 5 and the second leaves open those where it shows 1, 2, or 3. The conjunction leaves open a possibility when it is left open by both components; that is, it leaves open those where the die shows 1 or 3.

This is shown pictorially in Figure 2.1.2-1 below.

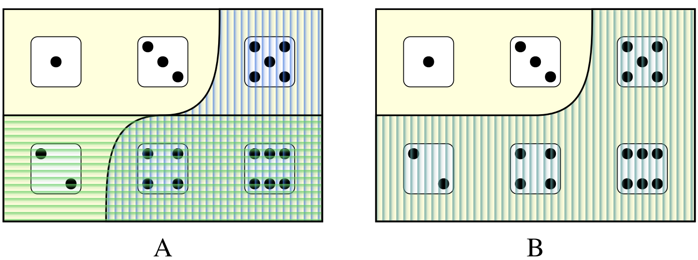

Fig. 2.1.2-1. Propositions expressed by two sentences (A) and their conjunction (B).

Here, each rectangle represents the space of all possible worlds. The die faces mark regions consisting of the possible worlds in which the die shows one or another number. In Figure 2.1.2-1A, the possibilities ruled out by the first component are at the bottom while those ruled out by the second component occupy the region at the right. The possibilities left open by the first component then form the region in the top half while those left open by the second are in the region at the left. Figure 2.1.2-1B shows the proposition expressed by the conjunction of these two sentences. The possibilities ruled out add up to form the shaded region; those left open are in the unhatched region at the top left where the ranges of possibilities left open by the original components overlap. These diagrams can be compared to the truth table for conjunction. The sort of worlds covered by first row of the table, worlds where both components are true, appear at the top left of the 2.1.2-1A; the other rows of the table correspond to the remaining three regions of the this diagram —those at the top right, the bottom left, and the bottom right, respectively.

Although the most fundamental approach to the deductive properties of the logical form will come through laws principles concerning its role as a conclusion or one among possibly many premises of an entailment, specific characteristics can often be highlighted most clear by its significance for relations between pairs of sentences, especially the positive relations of implication and equivalence. The following principles are some of the more important examples of this in the case of conjunction:

COMMUTATIVITY. *The order of conjuncts in a conjunction does not affect the content.* That is, $\varphi \wedge \psi \simeq \psi \wedge \varphi$.

ASSOCIATIVITY. *When a conjunction is a conjunct of a larger conjunction, the way components are grouped does not affect the content.* That is, $\varphi \wedge (\psi \wedge \chi) \simeq (\varphi \wedge \psi) \wedge \chi$.

IDEMPOTENCE. *Conjoining a sentence to itself does not change the content.* That is, $\varphi \wedge \varphi \simeq \varphi$.

COVARIANCE. *A conjunction implies the result of replacing a component with anything that component implies.* That is, if $\psi \vDash \chi$, then $\varphi \wedge \psi \vDash \varphi \wedge \chi$ and $\psi \wedge \varphi \vDash \chi \wedge \varphi$.

The names of these principles are terms used for analogous principles in other contexts. For example, you may have encountered the first two as names of principles for addition and multiplication, for both of which order and grouping do not matter. Conjunction shares the third property with numerical operations that produce the maximum or minimum of a pair of numbers, and this is not surprising since, if we think of truth values ordered so that falsity comes below truth, then the truth value of a conjunction is just the minimum of the truth values of its components.

The last property, covariance, says roughly that the content of a conjunction varies in the same direction as the content of its components. An analogous property holds for addition and the maximum and minimum operations (e.g., if $y \leq z$ then $\min(x, y) \leq \min(x, z)$) but not for multiplication when negative numbers are considered (e.g., $-2 \times 3 > -2 \times 4$ even though $3 \leq 4$). We cannot say that an increase or decrease in the content of one component will produce an increase or decrease, respectively, in the conjunction since information added or lost in a change to one component may be provided in any case by the other component. For example, although The sign had red letters on a blue background says more than does The sign had red letters, the conjunction The sign had red letters on a blue background, and the background was light blue is equivalent to The sign had red letters, and the background was light blue. This analogous to the fact that, $\min(2, 3) = \min(2, 4)$ even though $3 < 4$. What can be said is that, if the content of one component of conjunction increases, the content of the conjunction must increase if it changes at all.

One consequence of covariance is the following principle:

COMPOSITIONALITY. *Conjunctions are equivalent if their corresponding components are equivalent.* That is, if $\varphi \simeq \varphi'$ and $\psi \simeq \psi'$, then $\varphi \wedge \psi \simeq \varphi' \wedge \psi'$.

Although this follows from covariance (since equivalent components imply each other), it can hold when covariance does not and is so fundamental that, if conjunction did not satisfy it, we might hesitate even to count it as a logical form. Since sentences are logically equivalent when they express the same proposition, the principle says that conjunctions cannot express different propositions unless there is some difference in the propositions expressed by their components. Understanding the meanings of sentences to be the propositions expressed, the principle of compositionality tell us that the meaning of a conjunction is composed out of the meanings of its components in the particular way we label "conjunction."

Glen Helman 10 Aug 2010

2.1.3. Conjunction in English

Conjunction is most often marked by the word and, but there are English sentences without this word that also may be analyzed as conjunctions. First of all, there are quite a number of expressions—such as also, in addition, and moreover—that serve as stylistic variants of and. But conjunctions also may employ another group of words that are not simple stylistic variants of and. The principal example is the word but.

This may be a surprise. Although a sharp ear might detect a slight difference in meaning between and and moreover, the difference between and and but is unmistakable. Consider, for example, the following two sentences, which differ only in the use of these two words:

Adams spoke forcefully to the committee, and they agreed to the expenditure

Adams spoke forcefully to the committee, but they agreed to the expenditure.

These sentences would be used under different circumstances, and it may seem odd to count them as logically equivalent, which is what we must do if we are to analyze both as conjunctions of the same two components.

This is the first of several points at which we must recall the distinctions between truth and appropriateness and between implication and implicature. As was noted in 1.3.4, our concern is with only the first concept in each pair and thus with only certain aspects of meaning. Specifically, we count two sentences as equivalent if they have the same truth conditions. Any differences between their meanings that have no effect on their truth and falsity are irrelevant for our purposes.

So we must look more closely at the nature of the difference in meaning between and and but. It is clear that the second sentence above carries a suggestion of contrast between the two components—perhaps Adams spoke against the expenditure or the committee usually rejected Adams's advice—and it is also clear that the suggestion of contrast is absent in the first sentence. Now, suppose that the second sentence was used in a context where the suggested contrast is not present—perhaps the expenditure was approved because Adams spoke for it. The assertion of the second sentence would then be inappropriate, but would it be false?

Let us use the test of a yes-no question. Imagine that you attended a meeting were Adams persuaded a committee to agree to a certain expenditure and that later someone who had heard rumors of the proceedings asked you the question Is it true that Adams spoke forcefully to the committee, but they agreed to the expenditure?. How would you reply? This is something you must decide for yourself; but, for my own part, I would say something like, “Yes, but he spoke for the expenditure, not against it.” That is, I would give a yes-but answer, reacting to the sentence whose truth was asked about as one whose assertion would be true but inappropriate. And it is for this reason that I will suggest we analyze sentences formed using but and other similar words—such as however, though, and nonetheless—as conjunctions. These words are not just signs of conjunction; but their differences from and lie outside their effect on truth conditions.

There are cases of other sorts where analysis by conjunction is legitimate though not obvious. Sometimes, for example, there is no word at all marking the conjunction. The operation of conjunction produces a compound sentence that commits us to the truth of both its components, and there are linguistic devices other than the use of particular words that enable us to roll two claims up into one in this way. For example, the sentence It was a hot, windy day is equivalent to It was a hot and windy day and can be analyzed as the conjunction

It was a hot day $\wedge$ it was a windy day.

An analysis of a sentence might even separate a modifier from the expression it modifies. One common case of this is provided by adjectives used *attributively*—i.e., applied directly to the noun they modify. For example, we may treat Sam's car is a green Chevy as if it were Sam's car is a Chevy, and it's green. It is important to note that, for reasons discussed in the next section, these analyses work only because the adjectives appear in a predicate nominative employing the indefinite article—i.e., in the form represented by

X is a … Y

or by a similar form with a different tense. However, this is a very common pattern so there will be many occasions to apply this sort of analysis.

Another rather specific but important case of separating modifiers concerns relative clauses. There are really two cases here. The first is non-restrictive relative clauses—that is, ones marked off by commas. These can usually be analyzed as conjunctions. For example, Ann, who you met yesterday, called this morning can be understood as a conjunction of You met Ann yesterday and Ann called this morning.

The second sort of case is a restrictive relative clause—one not marked off by commas—appearing as part of a predicate nominative using the indefinite article. The grammatical pattern in this case is

X is a Y that …

or a similar pattern using a different tense or another relative pronoun (such as who or which). A sentence like this can be analyzed as a conjunction of X is a Y and the result of putting X in the expression marked by … at the point governed by the relative pronoun. For example, Sam's car is a Chevy that's green could be analyzed as the conjunction of Sam's car is a Chevy and Sam's car is green—i.e., analyzed in the same way as Sam's car is a green Chevy. But relative clauses of this sort can be used to express many sorts of modification other than the simple application of adjectives. One example is The speaker was a writer who Sam admired, which can be analyzed as the conjunction of The speaker was a writer and Sam admired the speaker; here the second conjunct has the subject of the original sentence as its direct object rather than its subject.

Glen Helman 03 Aug 2010

2.1.4. Limits on analysis

Although the presence of and, or another word used to mark conjunction, is a good sign that conjunction will be involved in a full analysis of a sentence, it does not mean that the sentence as a whole can be analyzed as a conjunction. For that to be possible, we must be able to identify components that are independent sentences, and there are a number of things that can keep us from doing that.

One thing that can interfere is the occurrence of indefinite articles and similar expressions. Consider the sentence A friend of Ann lives in Singapore and works in London. The claim it makes may well be true, but its truth would be at least mildly surprising. However, there would be much less surprise at the truth of a sentence analyzed as A friend of Ann lives in Singapore ∧ a friend of Ann works in London since there is no longer any implication that the same person does both. Of course, we could paraphrase the original sentence (a bit awkwardly) as A friend of Ann lives in Singapore and that person works in London, but that is of no help in analyzing it since the second clause relies on the first clause for the reference of the phrase that person and thus does not function as an independent sentence. Indeed, in spite of the occurrence of the word and, there is no way to analyze this sentence as a conjunction in which the references to Singapore and London appear in different components; its analysis must await our treatment of expressions involving the indefinite article in chapters 7 and 8. The indefinite article is one of a group of expressions also including some, every, and no that we will later study as *quantifier words*. Their presence will often preclude analysis of a sentence as a compound formed by a connective even though a word that ordinarily indicates that compound is present. Analysis as a compound formed by the connective is sometimes possible in such cases, but you should be wary if you find yourself being led to repeat a quantifier word when dividing the sentence into two components (as we would do by repeating a friend of Ann in the example above).

Similar problems can arise in other cases where we might expect to find a conjunction, as with attributive adjectives and relative clauses. For example, Tom forecast a hot and windy day next week is not equivalent to Tom forecast a hot day next week ∧ Tom forecast a windy day next week since the latter does not imply that the two forecasts are for the same day. This is the reason that 2.1.3 recommended such analyses only for predicate nominatives. In such cases, the implication that two adjectives are being applied to the same thing is insured by other aspects of the sentence, but you still need to be wary of duplicating other quantifier phrases—in, for example, the subject of the sentence—when you make the analysis. And this is true even for compound predicate adjectives: Sam's car was cheap and reliable is equivalent to Sam's car was cheap ∧ Sam's car was reliable but One model is cheap and reliable is not equivalent to One model is cheap ∧ one model is reliable.

Even when quantifier words are not involved, analyses by conjunction cannot always be used to separate modifiers from the words they modify. For example, it would be wrong to analyze Tristram is a large flea as Tristram is a flea ∧ Tristram is large because a sentence with this analysis entails that at least one flea is to be found among the large things of the world. The problem in this case is that an adjective modifying a noun has its meaning determined in part by the noun it is applied to; large indicates a different range of sizes when it is applied to fleas than when it is applied to elephants. This is an example of a phenomenon discussed in 1.3.6: vague terms have their meaning determined in part by their context of use. A noun can contribute to the context in which an adjective is used when the adjective is applied to the noun directly and also when the adjective follows the noun in a stream of discourse. This means that it also would be wrong to analyze Tristram is a flea and Tristram is large as Tristram is a flea ∧ Tristram is large, for the adjective large acquires part of its meaning from the noun flea in the English sentence. (But the way a noun affects the meaning of a vague adjective is not simple. Although the sentence No fleas are large speaks about fleas, the range of sizes indicated by large in this sentence is different from the range indicated by its use in Tristram is a flea and Tristram is large.)

But why does the same thing not happen with the conjunction Tristram is a flea ∧ Tristram is large? Although the symbol ∧ is closely related to the English conjunction and, it is not a simple abbreviation; and we do not assume that their contribution to the meaning of a sentence is exactly the same. The symbol ∧ (and the construction both … and … that we use as an alternative notation for it) are signs for the operation of conjunction. The conjunction of two sentences is a sentence that, in any context, has truth conditions that are related to those of its two components in the way shown by the table we considered earlier. And the stipulation that this is so *in any context* is a crucial one here; in particular, it need not be part of that context that either component has been asserted. So in the conjunction, we cannot assume that the meaning of the second component Tristram is large will be influenced by the meaning of the first component. In certain sorts of context, Tristram is large will have the same meaning as Tristram is large for a flea. But it is only in such contexts that Tristram is a flea ∧ Tristram is large has the same truth conditions as Tristram is a flea and Tristram is large, and our analyses should not depend on equivalences that hold only for certain contexts.

This indicates a further difference between our model of the operation of language and the way things work in English. Everything that is said in English has the potential of affecting the context of what follows it and, to a more limited extent, what precedes it. But when we analyze sentences, we treat their components as independent and as each understood in the same context. Our excuse for this limitation of our model is the same as that for many others: a model that was more accurate in this respect would require significant complications—and complications that no one yet understands very well.

Of course, we can analyze Tristram is a large flea as a conjunction after all if we modify the second component to remove its dependence on the context established by the assertion of the first. One way of doing that was suggested in passing above: we may use the conjunction Tristram is a flea ∧ Tristram is large for a flea. Here we have modified the second component to replace the implicit effect of the context with a more explicit indication of the range of sizes in question. Though generalizations about such matters are risky, something like this device can be applied in many cases where adjectives acquire part of their meaning from the surrounding context.

There are still other factors that can prevent the separation of attributive adjectives from the nouns they modify. We could be guilty of slander if we were to analyze Alfred is an alleged murderer as Alfred is a murderer ∧ Alfred is alleged to be a murderer. The difference between this and the example above is that the attributive adjective alleged modifies the meaning of a noun in a different way from an adjective like large. Adjectives like large narrow down the class of things marked out by the noun by adding a further property; in contrast, alleged shifts the membership of this class by adding as well as dropping members. The class of alleged murderers is not included in the class of murderers in the way the class of large fleas is included in the class of fleas. As a result, no analysis as a conjunction is possible.

While the issues of contextual dependence can also affect our ability to separate relative clauses from the nouns they modify, this latter problem does not occur for them. If we say Alfred is a murderer who is alleged to be one we already imply that Alfred is a murderer so analysis as a conjunction is possible. This means that one initial test for cases where we may separate an attributive adjective from the noun it modifies is to see if restatement using a relative clause changes the meaning. While That's an unknown Rembrandt is equivalent to That's a Rembrandt that is unknown and can be analyzed as a conjunction, That's a fake Rembrandt is not equivalent to That's a Rembrandt that is fake and cannot be analyzed in this way.

But, in the end, the test that an analysis must pass is that the conjunction we use to represent a sentence really has the same truth conditions. Since the truth table for conjunction is directly tied to the laws of entailment discussed in 2.1.1, one way to apply this test is to check whether the original sentence really entails both components of the analysis (when these are considered as independent sentences) and whether they, taken together, entail it. And we have used this test in the discussion of examples above; for example, because Alfred is an alleged murderer does not entail Alfred is a murderer, we cannot analyze the premise as conjunction with the conclusion as one of its conjuncts. Due to the problems associated with the contextual dependence of meaning, when applying this test, we must be careful not to fill out the meanings of terms in one of the sentences we compare by a surreptitious reference to another sentence.

2.1.5. Multiple conjunction

Although conjunction can compound sentences only two at a time, the word and in English can be used with any series of more than two items. To analyze a *serial conjunction* like He went to Gary, South Bend, and Fort Wayne, we need to regard the sentence as the result of two uses of conjunction, first to join two of the components and then to tack on a third. There are two ways of doing this and, although the associativity of conjunction noted in 2.1.2 tells that they have the same content, they arrive at their common meaning in different ways.

We can represent this difference in our symbolic notation by using parentheses:

He went to Gary ∧ (he went to South Bend ∧ he went to Fort Wayne)
(He went to Gary ∧ he went to South Bend) ∧ he went to Fort Wayne.

There are a number of ways of describing the difference displayed here. We can say, first, that in each case a different one of the two uses of conjunction is the *main* connective or the one at the *top level*. The main or top-level connective is the operation that would be used last in forming the sentence, and it marks the place the sentence would be broken first when it is decomposed. In the first sentence above, it is the first use of conjunction that is the main connective or the one at top level while, in the second sentence, it is the second use.

Another way of describing the difference between the two analyses is to speak of the *scope* of a connective, the part of the whole sentence that is made up of the connective and the components it applies to. Thus the scope of the first ∧ in the first of the sentences above is the whole sentence while the scope of the second ∧ is the portion in parentheses. This situation is reversed in the second sentence; there, the scope of the first ∧ is limited by the parentheses and is included in the scope of the second ∧.

He went to Gary ∧ (he went to South Bend ∧ he went to Fort Wayne)

(He went to Gary ∧ he went to South Bend) ∧ he went to Fort Wayne

So we say that the two examples differ in the *relative scope* of the two uses of conjunction. In one, the first use has *wider* scope; in the other, the second has wider scope.

These two ideas are depicted together in Figure 2.1.5-1. The main connective of each analysis appears quite literally at the top level, and the scope of each connective is the portion of the analysis that branches out from under it.

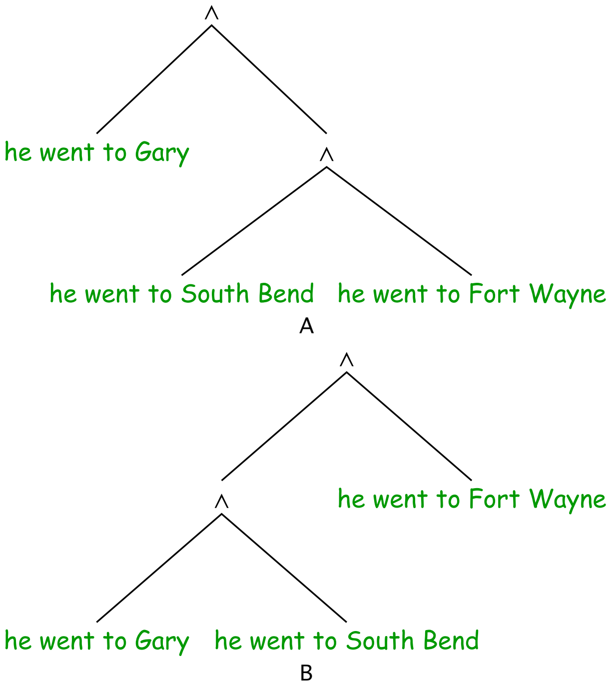

Fig. 2.1.5-1. Two analyses of a serial conjunction.

The parentheses of our symbolic notation for these analyses can be seen as a way of representing this sort of structure without resorting to two dimensions.

Although such scope distinctions make no difference in the truth conditions of English sentences, they can be marked syntactically—as in

He went to Gary and also South Bend and Fort Wayne
He went to Gary and South Bend and also Fort Wayne.

Here, also is used to emphasize the break made by one of the occurrences of and over the other. Use of punctuation is another way to emphasize one of the two ands and raise it to the top level—as in

He went to Gary—and South Bend and Fort Wayne
He went to Gary and South Bend—and Fort Wayne.

In the absence of devices like the use of also, syntactic grouping, or punctuation, the normal order of reading probably would lead us to interpret the second and as the last one used in forming the sentence—that is, as the main operation.

Still another common way of making such distinctions is to exploit the power of and to conjoin words or phrases as well as complete clauses. For example, compare (in which the conjoined words and phrases are underlined)

He went to Gary and to South Bend and Fort Wayne

He went to Gary and South Bend and to Fort Wayne

In each case, one and conjoins prepositional phrases and the other conjoins nouns to form the object of one of these phrases.

A final way of representing scope distinctions in English is one we have adapted to represent conjunction using the expression both … and …. All things being equal, we will interpret the second of two ands as having the wider scope, but this presumption can be defeated by adding the word both to get, for example

He went to Gary and both South Bend and Fort Wayne.

This sentence has the form $\varphi \land (\psi \land \chi)$; and, in general, the word both has roughly the same effect as a left parenthesis in our symbolic notation. Indeed, when we represent the forms we have been considering using English notation we get this:

both he went to Gary and both he went to South Bend and he went to Fort Wayne

both both he went to Gary and he went to South Bend and he went to Fort Wayne.

The word both appears here just where a left parenthesis would in a symbolic analysis if we were to add parentheses surrounding the whole sentence.

(He went to Gary $\land$ (he went to South Bend $\land$ he went to Fort Wayne))

((He went to Gary $\land$ he went to South Bend) $\land$ he went to Fort Wayne)

Of course, our English notation uses both in many cases where both would not appear in ordinary English and even where a left parenthesis would not ordinarily appear in our symbolic notation. This is because the English notation is designed to make the scope of connectives unambiguous in cases where ordinary English is ambiguous. And without anything to serve as a corresponding right parenthesis, the word both may be needed in some places where a left parenthesis is not.

For example, suppose we attempt to express the form $(\varphi \land \psi) \land \chi$, using and in place of $\land$ and both in place of the left parenthesis. We would get

both φ and ψ and χ.

If we take both to mark the left end of the scope of the first conjunction (as the left parenthesis does in the symbolic expression), we are still left with no indication of the right end of its scope: is the second component only ψ, or is it the whole of ψ and χ? If we supply a both for every and, we can write $(\varphi \wedge \psi) \wedge \chi$ as

both both φ and ψ and χ.

This is hardly elegant prose, but it does make the grouping definite; finding a second both immediately following the first, we know the first component of the main conjunction is itself a conjunction. Of course, we could also mark scope using parentheses. It may seem odd to do this if we are using English notation, too; but it is possible to mix the two forms, and it can sometimes be helpful to indicate a logical form by combining the word and with grouping marked by parentheses.

Although scope distinctions can be made in English in these ways, the English and is often applied to a series of items that are all on the same level. It would be possible to treat conjunction as an operation that was similar to the English and in this respect, but it would cost us the trouble of more complex accounts of the properties of conjunction without yielding much greater insight. We can (and often will) mimic the way addition and multiplication are usually treated in algebra and drop parentheses when they make no difference in the value of an expression. This introduces no real complications but it has limitations. Since our principles concerning conjunction will be stated only for 2-component conjunctions, we can apply them to a *run-on conjunction* like $\varphi \wedge \psi \wedge \chi$—or, in English notation, φ and ψ and χ—only after we have chosen one of the two conjunction symbols as marking the main operation. And, although we could regard either the first or last of the three components as a component of the top-level conjunction, the middle one ψ always ends up as a component of the lower level conjunction, so we really have not put the three components on the same level.

2.1.6. Some sample analyses

Here are a few example analyses written out in full as models for the exercises to this section. In each case a few comments follow the actual analysis.

Roses are red and violets are blue
Roses are red $\land$ violets are blue

R $\land$ B
both R and B

R: roses are red; B: violets are blue

As a last step here, unanalyzed have been abbreviated components with capital letters in order to highlight logical forms. The final form is stated both symbolically and using English notation, something that will be done also in the examples to follow.

The next example is worked out in two steps, first analyzing the whole sentence as a conjunction and then analyzing one of its components.

It's cool even though it's bright and sunny
It's cool $\land$ it's bright and sunny
It's cool $\land$ (it's bright $\land$ it's sunny)

C $\land$ (B $\land$ S)
both C and both B and S

C: it's cool; B: it's bright; S: it's sunny

The parentheses in the final result correspond to the grouping of bright and sunny together in the predicate of the second clause of the original sentence.

In the following example, it would not be wrong to use parentheses (or grouping with both), but that would be an artifact of our analysis and correspond to nothing in the English.

He was cool, calm, and collected
He was cool $\land$ he was calm $\land$ he was collected

C $\land$ M $\land$ T
C and M and T

C: he was cool; M: he was calm; T: he was collected

Accordingly, the analysis uses a run-on conjunction in the symbolic version, and use of both is similarly suppressed in the English statement of the form. If grouping were used here, either way conjunction might be assigned widest scope.

Finally, there can be cases where some grouping reflects the structure of the English, but other grouping does not.

It is a two-story brick building with a slate roof
It is a two-story brick building $\wedge$ it has a slate roof
(it is a building $\wedge$ it is made of brick $\wedge$ it has two stories) $\wedge$ it has a slate roof

$(B \wedge R \wedge T) \wedge S$
(B and R and T) and S

B: it is a building; R: it is made of brick; S: it has a slate roof; T: it has two stories

No grouping is used within the first three components because it is not obvious that any is imposed by the phrase two-story brick building. The English notation employs parentheses because there is no good way of indicating the combination of run-on conjunction with ordinary conjunction using both.

As in the last example, there would be nothing wrong with imposing a grouping here. If we were to group the first three components to the left, we would end up with the following in symbols and English:

$((B \wedge R) \wedge T) \wedge S$
both both both B and R and T and S

In the English notation, each of the boths tells us that a certain component is a conjunction—first the whole sentence, then its first component, and finally the first component of this component—and this settles the scope of the ands that follow.

The value of English notation does not lie in the possibility of making such a calculation but rather in our ability to understand the significance both automatically; however, that ability is limited to fairly simple forms, and a row of three boths is hard to follow without reflection. (To cite a standard example of a similar limitation in the case of a different sort of grouping, it is just possible to understand Bears bears fight fight to say what is said by Bears that bears fight fight—i.e., so that the first bears is modified by a relative clause bears fight and is the subject of the second fight; but it is virtually impossible to understand Bears bears bears fight fight fight as anything other than a cheer, even though it is grammatically possible for it to say something that might be expressed by Bears (which bears (that bears fight) fight) fight.)

2.1.7. Logical forms

We will conclude this first look at analysis by considering its results in more general terms. The aim of analysis is to uncover logical form. While it is natural to speak of the result of an analysis as *the* logical form of the sentence that was analyzed, a sentence will usually have many logical forms of differing complexity. Many of these may be displayed as we carry out an analysis step by step. Consider, for example, the following analysis of a fairly complex sentence:

He went to Gary, South Bend, and Fort Wayne, leaving at dawn and returning after dark

He went to Gary, South Bend, and Fort Wayne ∧ he left at dawn and returned after dark

(he went to Gary and South Bend ∧ he went to Fort Wayne) ∧ he left at dawn and returned after dark

((he went to Gary ∧ he went to South Bend) ∧ he went to Fort Wayne) ∧ he left at dawn and returned after dark

((he went to Gary ∧ he went to South Bend) ∧ he went to Fort Wayne) ∧ (he left at dawn ∧ he returned after dark)

((G ∧ S) ∧ F) ∧ (L ∧ R)

both both both G and S and F and both L and R

F: he went to Fort Wayne; G: he went to Gary; L: he left at dawn; R: he returned after dark; S: he went to South Bend

The first line exhibits the sentence without further analysis, the second shows it as a conjunction, the third as a conjunction whose first component is a conjunction, and so on. (The first component might have been analyzed as a run-on conjunction; but, for the purposes of this example, we need a fully specified structure.)

Each line ascribes a form to the sentence, and if we ignore the identity of unanalyzed components, this is a form that the sentence shares with many other sentences. These abstract forms are indicated below (in the order in which they appear in the analysis) with symbolic notation on the left and a description of the form on the right:

φ	sentence
$\psi \wedge \chi$	conjunction
$(\zeta \wedge \xi) \wedge \chi$	conjunction of (i) a conjunction and (ii) a sentence
$((\mu \wedge \nu) \wedge \xi) \wedge \chi$	conjunction of (i) a conjunction whose first component is a conjunction and (ii) a sentence
$((\mu \wedge \nu) \wedge \xi) \wedge (\theta \wedge \upsilon)$	conjunction of (i) a conjunction whose first component is a conjunction and (ii) a conjunction

The sentence has still further forms that might have appeared in the course of our analysis if we had reached the final result in a different way. One example is $\psi \wedge (\theta \wedge \upsilon)$, a conjunction of (i) a sentence and (ii) a conjunction.

It is important to recognize all the different forms a sentence has, even those that correspond to very partial analyses of it. Each represents a class of sentences that may share important logical properties with the sentence we are focusing on. For example, the sentence above will share some of its logical properties with all sentences, others with all conjunctions, still others with conjunctions whose first components are conjunctions, and so on.

We will apply the term *component* to any sentence that appears on any level of analysis of a given sentence. In particular, a sentence is a component of itself. We will distinguish those components of a compound to which the main connective applies as the *immediate* components of the compound, and we will refer to those that appear unanalyzed at the last stage of an analysis as the *ultimate* components (on that analysis). We will often refer to the ultimate components of a sentence also as *unanalyzed*. In the example above, the immediate components of the initial sentence are the two sentences separated at the second line of the analysis; and the ultimate components are those abbreviated with capitals at the end. Although, in principle, both roman capital letters and the lower case Greek letters may stand for any sentences, in practice, we will reserve capital letters for sentences we do not analyze further. Such sentences are ultimate components of themselves and of any larger compounds in which they appear.

2.1.8. Interpretations

In passing from a sentence to any of its logical forms, we abstract from the specific components that we replace by variables. In general, we also abstract from the proposition expressed by the sentence and from its truth value. Except in special cases, such as forms that are shared only by tautologies, a logical form does not express a proposition or have a truth value, but we may introduce such semantic features by *interpreting* the form.

We will consider two sorts of interpretation, an *extensional interpretation*, that provides a truth value only, and an *intensional interpretation*, which provides the proposition expressed and thus a truth value not only for the actual world but for every possible world. These two sorts of interpretation will be used for different purposes, so it will usually be clear from the context which sort is relevant; and, when this is clear, we will use the term *interpretation* without qualification.

The term *intensional* (spelled with an *s*) and the term *extensional* derive from a traditional distinction between, on the one hand, the means by which a term picks out a class of objects and, on the other, the class of objects it picks out. Terms that pick out the same class of objects in different ways have the same *extension* but different *intensions*. For example, if the population of Crawfordsville is 14287, the terms city with a population greater than 14287 and city more populous than Crawfordsville have the same extension but different intensions. One way to see that the two terms have different intensions is to notice that they would pick out different classes of cities if the population of Crawfordsville were not 14287.

During the past century, the concepts of intension and extension have been extended to terms that pick out single objects rather than classes of objects, so we can say that the definite descriptions the author of *Poor Richard's Almanack* and the inventor of the lightning rod both have Benjamin Franklin as their extension though they differ in their intensions.

The distinction between the object a term refers to and the way the term refers to this object is sufficiently analogous to the distinction between the truth value of a sentence and the proposition it expresses that the concepts of intension and extension are now also applied to sentences. So Indianapolis is the capital of Indiana and Sringfield is the capital of Illinois could be said to have the same extension (i.e., the value **T**) but to have different intensions. In general, the extension of sentence is the sentence's truth value while the intension is the proposition that the sentence expresses.

Since the only general way we have to specify propositions is by using sentences that express them, intensional interpretations will be specified by assigning sentences to variables. (This assumes were are working with a fixed context of use, so sentences express propositions.) This assignment is the exact inverse of the process of abbreviating ultimate components by capital letters, and we will use the same notation for the association of letters and sentences in both. For example, we can give an intensional interpretation of the form (A ∧ B) ∧ C by making the following assignment of sentences to the variables that mark its ultimate components.

A: I got it apart;
B: I don't know how I got it apart;
C: I couldn't get it together again

Since the sentences assigned to variables serve only to specify propositions, we will not be concerned about their logical forms; they may be as simple or complex as we wish.

Especially in later chapters, the proposition assigned to a compound sentence by an intensional interpretation may not be apparent until we find an idiomatic English sentence that expresses the same proposition. This can be done by a step-by-step process of *synthesizing* English that reverses the process of analysis. For the example above, this might proceed as follows:

(I got it apart ∧ I don't know how I got it apart) ∧ I couldn't get it together again
I got it apart but I don't know how ∧ I couldn't get it together again
I got it apart but I don't know how, and I couldn't get it together again

Of course, other wording is possible here, and the process of synthesizing English will rarely have a unique correct result.

Extensional interpretations are easier to manage and will often provide all the information we need. We will adapt the tabular notation used for truth tables.

A	B	C	(A ∧ C)	∧	(B ∧ C)
T	F	T	T	Ⓕ	F

Variables are listed at the left with the assigned value under each of them. The whole form we are interested in is displayed to their right. The values of the larger components may be calculated by using the truth table for conjunction just as a multiplication table may be used to calculate the numerical value of a product: we find the values of the smallest components first and use these to calculate the values of larger components. The truth value calculated for each compound component is displayed below the main connective of that component. The value for the sentence as a whole is shown circled. Our interest will generally be only in this final value, but examples in this text will usually also display the intermediate values in order to show how the final value was reached.

2.1.s. Summary

1 The prime role of the logical word and is to mark the use of a connective, called conjunction, that serves to form a compound sentence (also called a conjunction) from component sentences that may be referred to as its conjuncts. The process of interpreting a sentence as a conjunction is analysis. We use the sign $\wedge$ (logical and) as symbolic notation for the operation of conjunction, marking the scope of a conjunction by parentheses. Alternatively, we can write a conjunction $\varphi \wedge \psi$ as both φ and ψ, where both plays the role of a left parenthesis. The two forms can be mixed using and to mark conjunction and parentheses to mark scope. We will use capital letters to stand for unanalyzed components as we use lower case Greek to stand for any sentences, analyzed or not.

2 The effect of conjunction on the truth conditions of the compounds formed using it may be described in a truth table showing the compound to be true if and only if both components are true. The truth table specifies a truth function, so conjunction can be said to have a truth function as its meaning. Some of the properties conjunction has in virtue of its meaning have standard names. It is commutative, associative, and idempotent (i.e., the order, grouping, and number of conjuncts does not affect the content of a sentence formed using conjunction, perhaps repeatedly); and it is covariant (adding or reducing the content of a component makes the content of the conjunction vary in an analogous way).

φ	ψ	$\varphi \wedge \psi$
T	T	T
T	F	F
F	T	F
F	F	F

3 Conjunction is marked in English by stylistic variants of and as well as by but and similar words. Conjunction also can appear without explicit indication, particularly through the use of modifiers like attributive adjectives and relative clauses.

4 Care is needed to be sure that such modifications can be captured by conjunction and to identify components that make independent contributions to the compound. The presence of quantifier words can preclude analysis as a conjunction even when the word and is present.

5 Since conjunction is used to combine only two components, uses of conjunction to combine more than two in a multiple conjunction will involve two or more connectives of differing scope, the one with widest scope counting as the main connective of the sentence. Such differences in scope can be marked in several ways in English but such

markings may be absent in a serial conjunction. Some of the effect of serial conjunction without scope distinctions can be achieved by run-on conjunctions, such as $\varphi \wedge \psi \wedge \chi$, which suppress parentheses.

6 In all but the simplest cases, the analysis of conjunctions will find components that are themselves conjunctions. The result of an analysis will exhibit this structure using symbolic and English notation. Although it is never wrong to mark the scope of conjunction within serial conjunctions, the resulting differences in the scopes of connectives are more significant in some cases than in others.

7 The analysis of the logical form of a sentence can occur in stages in which we identify the immediate components of a compound, any immediate components of these, and so on. The last components arrived at are the ultimate components of the analysis; the full class of components includes them as well as all other sentences that could appear in the course of analysis (including the analyzed sentence itself). A sentence will usually have many logical forms representing different partial analyses of it.

8 We can specify a proposition or a truth value for a logical form by means of an intensional or extensional interpretation, assigning truth values or sentences, respectively, to its ultimate components. A sentence expressing the proposition provided by an intensional interpretation can be found by carrying out a process of synthesis that reverses the process of analysis. The truth value provided by an extensional interpretation can be found by calculation using the truth table for conjunction. The tabular notation used to write the truth table of conjunction may be used also to describe extensional interpretations and the values that they give to compound forms.

2.1.x. Exercises

1. Analyze each of the following sentences in as much detail as possible.

a. Mike visited both London and Paris.

b. Ann wanted white wine but Bill and Carol wanted red.

c. It will rain and clear off, but it will rain.

d. That is a new but growing market.

e. Confucius is affable but dignified, austere but not harsh, polite but completely at ease. (*Analects* 7:37)

f. Although Tim lost his glasses and his wallet, each was returned.

g. Tim lost his glasses and his wallet, and one person found both.

2. Restate each of the following forms, putting English notation into symbols and vice versa (e.g., both A and B becomes $A \land B$, and $A \land B$ becomes both A and B). Indicate the scope of connectives in the result by underlining.

a. both A and both B and C

b. both both A and B and C

c. $(A \land B) \land (C \land D)$

d. $A \land ((B \land C) \land D)$

e. $(A \land (B \land C)) \land D$

f. both both both A and B and C and D

3. The logical forms below are followed by intensional interpretations of their unanalyzed components. In each case, synthesize an idiomatic English sentence that expresses the corresponding interpretation of whole form. Remember that there may be more than one correct answer.

a. $(V \land F) \land R$

[F: Fred visited Florence; R: Fred spent a week in Rome; V: Fred visited Venice]

b. $(J \land (S \land F)) \land K$

[F: he was fair; J: he was a judge; K: he had an excellent knowledge of the law; S: he was stern]

c. $(C \land T \land H) \land (W \land F \land S)$

[C: we arrived cold; F: we left stuffed; H: we arrived hungry; S: we left sleepy; T: we arrived tired; W: we left warm]

d. $O \land O$

[O: Old King Cole was a merry old soul]

4. Calculate truth values for all compound components of the forms below using the extensional interpretation provided in each case.

a.

A	B	C	$A \land (B \land C)$
T	T	F	

b.

A	B	C	D	$((A \land D) \land C) \land (B \land A)$
T	T	F	T	

2.1.xa. Exercise answers

1. **a.** Mike visited London ∧ Mike visited Paris

L ∧ P

both L and P

L: Mike visited London; P: Mike visited Paris

b. Ann wanted white wine ∧ Bill and Carol wanted red wine

Ann wanted white wine ∧ (Bill wanted red wine ∧ Carol wanted red wine)

A ∧ (B ∧ C)

both A and both B and C

A: Ann wanted white wine; B: Bill wanted red wine; C: Carol wanted red wine

c. It will rain and clear off ∧ it will rain

(it will rain ∧ it will clear off) ∧ it will rain

(R ∧ C) ∧ R

both both R and C and R

C: it will clear off; R: it will rain

d. That is a market ∧ that is new relative to other markets but growing

That is a market ∧ (that is new relative to other markets ∧ that is growing)

M ∧ (N ∧ G)

both M and both N and G

G: that is growing; M: that is a market; N: that is new relative to other markets

e. Confucius is affable but dignified ∧ Confucius is austere but not harsh ∧ Confucius is polite but completely at ease

(Confucius is affable ∧ Confucius is dignified) ∧ (Confucius is austere ∧ Confucius is not harsh) ∧ (Confucius is polite ∧ Confucius is completely at ease)

(A ∧ D) ∧ (S ∧ H) ∧ (P ∧ E)

(both A and D) and (both S and H) and (both P and E)

A: Confucius is affable; D: Confucius is dignified; E: Confucius is completely at ease; H: Confucius is not harsh; P: Confucius is polite; S: Confucius is austere

f. Tim lost his glasses and his wallet ∧ Tim's glasses and wallet were each returned

(Tim lost his glasses ∧ Tim lost his wallet) ∧ (Tim's glasses were returned ∧ Tim's wallet was returned)

$$(G \wedge W) \wedge (R \wedge T)$$

both both G and W and both R and T

G: Tim lost his glasses; R: Tim's glasses were returned; T: Tim's wallet was returned; W: Tim lost his wallet

g. Tim lost his glasses and his wallet ∧ one person found both Tim's glasses and his wallet

(Tim lost his glasses ∧ Tim lost his wallet) ∧ one person found both Tim's glasses and his wallet

$$(G \wedge W) \wedge O$$

both both G and W and O

G: Tim lost his glasses; O: one person found both Tim's glasses and his wallet; W: Tim lost his wallet

Note: One person found both Tim's glasses and his wallet cannot be analyzed further because One person found Tim's glasses ∧ one person found Tim's wallet does not imply that the same person found both.

2. **a.** $A \wedge (B \wedge C)$

b. $(A \wedge B) \wedge C$

c. both both A and B and both C and D

d. both A and both both B and C and D

e. both both A and both B and C and D

f. $((A \wedge B) \wedge C) \wedge D$

3. **a.** (Fred visited Venice ∧ Fred visited Florence) ∧ Fred spent a week in Rome

Fred visited Venice and Florence ∧ Fred spent a week in Rome

Fred visited Venice and Florence, and he spent a week in Rome

b. (he was a judge ∧ (he was stern ∧ he was fair)) ∧ he had an excellent knowledge of the law

(he was a judge ∧ he was stern but fair) ∧ he had an excellent knowledge of the law

He was a stern but fair judge who had an excellent knowledge of the law

c. (we arrived cold ∧ we arrived tired ∧ we arrived hungry) ∧ (we left warm ∧ we left stuffed ∧ we left sleepy)

We arrived cold, tired, and hungry ∧ we left warm, stuffed, and sleepy

We arrived cold, tired, and hungry; and we left warm, stuffed, and sleepy

d. Old King Cole was a merry old soul ∧ Old King Cole was a merry old soul

Old King Cole was a merry old soul, and a merry old soul was he

4. Numbers below the tables indicate the order in which values were computed

a.

A	B	C	A	∧	(B	∧	C)
T	T	F		F		F	
				2		1	

b.

A	B	C	D	((A	∧	D)	∧	C)	∧	(B	∧	A)
T	T	F	T		T		F		F		T	
					1		2		3		1	

Glen Helman 05 Aug 2010

2.2. Proofs: analyzing entailment

2.2.0. Overview

We can get some insight into deductive logic by looking at basic principles of entailment, but more will come by looking at how these principles may be combined in proofs.

2.2.1. Proofs as trees

The simplest way of combining deductive principles takes the shape of a tree in which premises, premises from which these premises are concluded, and so on, grow and branch from the final conclusion.

2.2.2. Derivations

Although writing a proof as a tree can make its structure very explicit, we will mainly use a compact notation that more closely matches the patterns that are used when deductive reasoning is put into words.

2.2.3. Rules for derivations

In the context of derivations, principles of entailment take the form of rules that direct the search for a proof.

2.2.4. An example

All derivations that involve conjunction alone share many features; we will look closely at one typical example.

2.2.5. Two perspectives on derivations

Derivations have aspects that reflect both tree-form and sequent proofs; the latter aspect will prove especially important.

2.2.6. More rules

Tautology and absurdity provide a first example of derivation rules for logical forms other than conjunction.

2.2.7. Resources

In order to plot a course in constructing a proof for a given conclusion, we need to keep track of not only the premises but also the conclusions that have already been reached.

2.2.1. Proofs as trees

Our study of entailments involving conjunction will rest on the principles discussed in 2.1.1. These are shown below, in symbolic form on the left and in English on the right:

$\varphi \wedge \psi \vDash \varphi$	both φ and $\psi \vDash \varphi$
$\varphi \wedge \psi \vDash \psi$	both φ and $\psi \vDash \psi$
$\varphi, \psi \vDash \varphi \wedge \psi$	$\varphi, \psi \vDash$ both φ and ψ.

We will refer to the first two of these patterns as *extraction* (*left* and *right* extraction to distinguish them) and to the third simply as *conjunction*. To establish particular cases of entailment, we will want to put together instances of these general patterns and, eventually, instances of other patterns, too.

What may be the most direct notation for doing that employs something like the two-dimensional form we have used for arguments, with the conclusion below the premises and marked off from them by a horizontal line. In order to make the premises of a multi-premised argument available to serve as conclusions of further argument, we will spread them out horizontally. In this style of notation, the basic patterns for conjunction take the following forms (where abbreviations of their names are used as labels):

$$\text{Ext}\,\dfrac{\varphi \wedge \psi}{\varphi} \qquad\qquad \text{Ext}\,\dfrac{\varphi \wedge \psi}{\psi} \qquad\qquad \text{Cnj}\,\dfrac{\varphi \qquad \psi}{\varphi \wedge \psi}$$

Arguments exhibiting these patterns can be linked by treating the premises of one argument as conclusions of other arguments. For example, the following shows that $(A \wedge B) \wedge C$ is a valid conclusion from the two premises A and $B \wedge C$:

$$\text{Cnj}\,\dfrac{\text{Cnj}\,\dfrac{A \qquad \text{Ext}\,\dfrac{B \wedge C}{B}}{A \wedge B} \qquad \text{Ext}\,\dfrac{B \wedge C}{C}}{(A \wedge B) \wedge C}$$

The ability to put entailments together in this way rests on the general laws of entailment discussed in 1.4.6. The law for premises enables us to begin; it shows that the premises A and $B \wedge C$ entail the tips of the branches of this tree-like proof. Repeated uses of the chain law then enable us to add conclusions drawn using the principles for conjunction, and we work our way down the tree showing that the original set of premises entails each intermediate conclusion and, eventually, $(A \wedge B) \wedge C$. For example, just before the end, we know that our original premises entail each of the premises of the final conclusion—i.e., that A, B $\wedge$ C $\vDash$ A $\wedge$ B and A, B $\wedge$ C $\vDash$ C. The chain law then enables us to combine these entailments with the fact that A $\wedge$ B, C $\vDash$ (A $\wedge$ B) $\wedge$ C (a case of Conjunction) to show that A, B $\wedge$ C $\vDash$ (A $\wedge$ B) $\wedge$ C.

The simplicity of these tree-form proofs makes them useful for studying the general properties of proofs, but actually writing them out can become awkward. In the next section, we will look at a different sort of notation that makes it easier to write out proofs. It is most closely tied to a different way of stating the basic principles for conjunction that builds in the use of the chain law described above. Rather than pointing to particular valid arguments involving conjunction, these principles describe general conditions under which any arguments involving conjunction are valid.

LAW FOR CONJUNCTION AS A PREMISE. $\Gamma, \varphi \wedge \psi \vDash \chi$ if and only if $\Gamma, \varphi, \psi \vDash \chi$

LAW FOR CONJUNCTION AS A CONCLUSION. $\Gamma \vDash \varphi \wedge \psi$ if and only if both $\Gamma \vDash \varphi$ and $\Gamma \vDash \psi$

These principles can be seen to hold by the comparing the sort of possible worlds each side of the if and only if rules out.

The if parts of these principles reflect the validity of arguments of the forms Ext and Cnj, respectively, together with the chain law. The only if parts tell us that the validity of the arguments on their left sides can always be established in this way. For example, the only if part of the second tells us that, if a conjunction is a valid conclusion, then the premises needed to reach it by Cnj are bound to be valid conclusions also; so it should be possible to establish what we need to in order to apply Cnj.

When conjunction is the only connective employed in our analysis of sentences, applying these two principles repeatedly will eventually bring us back to arguments whose premises and conclusions are all unanalyzed components. An argument like that will be valid if its conclusion is among its premises; and, if the unanalyzed components making it up are logically independent, that is the only way it can be valid. This means that the two principles for conjunction combine with the law for premises to provide a complete account of validity for arguments involving only conjunction.

These two principles could be used to show that A, B $\wedge$ C $\vDash$ (A $\wedge$ B) $\wedge$ C—that is, to show what we showed in the earlier tree-form proof—in the following way:

$$\text{conj. as prem.}\;\dfrac{\text{conj. as concl.}\;\dfrac{\text{conj. as concl.}\;\dfrac{\text{law for prems.}\;\dfrac{}{A, B, C \vDash A} \qquad \text{law for prems.}\;\dfrac{}{A, B, C \vDash B}}{A, B, C \vDash A \land B} \qquad \text{law for prems.}\;\dfrac{}{A, B, C \vDash C}}{A, B, C \vDash (A \land B) \land C}}{A, B \land C \vDash (A \land B) \land C}$$

Like the tree-form proofs, this second way of writing proofs is being used only temporarily, but it is useful to have a name for it. It is close in form to a standard notation for proofs in which the separate claims of entailment are called *sequents*, so we will refer to proofs of this sort as *sequent proofs*.

A sequent proof can be read top to bottom as a tree-form proof like the earlier one except for two differences: (i) we are now reasoning about claims of entailment rather than unspecified sentences, and (ii) we are using principles of entailment rather than valid patterns of argument that apply to sentences of any sort. The examples we have looked have the further difference that the sequent proof begins with no assumptions—so the horizontal lines at the top have nothing above them—since the premises for principles at the next level down are provided by the law for premises, and that principle states categorically that certain arguments are valid (rather than making a conditional if and only if claim).

A sequent proof can also be read from the bottom up—that is, it can be understood to grow like a tree with a claim of entailment at its root. Looking at in this way, a sequent proof serves to investigate the conditions under which the entailment at its root holds. Or, more pointedly, it serves to search for ways in which that entailment might fail, ways of dividing its premises from its conclusion. In the example above, that search ends at the tips of branches when we run into arguments whose conclusions are among their premises.

Any notation for writing out proofs is more than we need to settle questions of entailment involving only conjunction. But the complications introduced by the logical forms we will consider in later chapters make it useful to have some system of notation, and, because we have simpler ways of seeing that entailments hold in the case of conjunction, it will be easier to see how this notation works if we develop it now. Neither the tree-form proofs nor the sequent proofs are the sort of notation we will actually adopt; but that more compact notation will have ties to both of them, and it will useful to look at them from time to time since they do exhibit quite clearly some features of the compact notation that are disguised by its compactness.

2.2.2. Derivations

Both ways of writing proofs that we considered in the last section involved trees that spread horizontally. The more compact notation that we will actually use will be more linear, though still somewhat two-dimensional. We will gain compactness by listing premises and conclusions in a more-or-less vertical way and by minimizing the repetition of premises that are used draw a number of conclusions. We will still need a tree structure to keep track of the premises relevant any given point, but this will involve rather stunted trees that grow horizontally from left to right.

Compactness is not all we will gain with this notation. It is designed to incorporate more directly the process of proof discovery, and it will approximate the ways proofs are normally stated in language. Indeed, although we will not emphasize this aspect of it, the notation for proofs could be thought of as a notation for analyzing the form of proofs presented in English that is in some respects analogous to our symbolic notation for analyzing the logical forms of sentences.

The system to be developed here falls into a broad class often referred to as *natural deduction systems* because they replicate, to some extent, natural patterns of reasoning. Such systems were first set out in full in the 1930s by G. Gentzen and also by S. Jaskowski, but some of the key ideas can be found already in the Stoic philosopher Chrysippus (who lived in the 3rd century BCE). The notation we will be using is an adaptation of notation introduced by F. B. Fitch but our approach to these systems will be influenced heavily by the "semantic tableaux" of E. Beth. (Their ideas are now a little over 50 years old.)

This system, which we will call a *system of derivations*, will employ a perspective on proofs that we adopted in the last section whenever we considered ways of restating claims of entailment. If we ask whether an entailment holds, we find ourselves faced with the task of reaching the conclusion from the premises (or showing that it cannot be reached). Let us think of the conclusion as our *goal* and of the premises as the *resources* we have available in trying to reach that goal. Until we reach the goal, it is separated from our resources by a *gap* that it is our aim to close.

We begin in the state shown in Figure 2.2.2-1, with a single gap between the premises and conclusion the argument whose validity we are trying to establish.

premise
premise ← resources
premise

← gap

conclusion ← goal

Fig. 2.2.2-1. The initial state of a derivation.

The premises of the argument (if it has any) are written above a horizontal line, and the conclusion is written below a second line. The space in between the horizontal lines marks the gap and will be filled in with additional resources and new goals as the derivation develops. (The vertical line on the left will be discussed later.)

We will approach the problem of closing the initial gap (or showing that it cannot be closed) step by step. At each step, either we will plan the way a goal may be reached or we will exploit resources, usually by drawing one or more conclusions from them. In making a step of either sort, we will restate our problem with different goals or resources, and we will say that, by this restatement, we are *developing* the derivation. When it is seen from this perspective, the problem of closing a gap is a problem of connecting available premises with desired conclusions. In developing a derivation, we work forward from premises and backward from conclusions in hopes of making this connection.

Either process may lead us to divide a gap in two. In the case of conjunction, this will happen when we plan to reach a goal $\varphi \land \psi$ by first concluding φ and ψ separately, for we will then set φ and ψ as separate preliminary goals and there will be a gap before each of them. This development of our initial problem by restating it and perhaps dividing it into subproblems will be expressed in a sort of tree structure. However, a derivation will be written as a more or less vertical list of sentences. The subgoals that we plan to reach in order to go on to a further goal will be written one above the other, each preceded by space for further growth, and conclusions we reach by exploiting resources will be written in at the top of a gap. In order to indicate the tree structure of problems and subproblems within this vertical list of sentences, we will need to mark up the derivation in various ways.

We will employ two main devices for doing this. One is the numbering of stages and sentences added at those stages. The other device is a system of vertical lines like the line at the left in Figure 2.2.2-1. These lines will be called *scope lines*, and they will serve us in a number of ways. First of all, new scope lines will be introduced as we analyze goals, with a separate scope line serving to mark the portion of the derivation devoted to each subgoal. The scope line will indicate the portion of derivation where a given subgoal is the goal we are aiming at, and it is in this sense that the scope line marks scope of the subgoal. As scope lines accumulate, they will be nested, some to the right of others, in a way that indicates the tree structure of the proofs. In later chapters, proofs will sometimes involve assumptions beyond the initial premises, and scope lines will then also serve to mark the portions of a proof in which these assumptions are operative—that is, they will serve to mark the scope of assumptions as well as goals. Later still, the scope lines will be labeled to indicate vocabulary that has a special role in the portion of a derivation marked by the scope line.

At any stage in the development of a derivation, each gap will have certain *active resources*. These are resources available for use in the gap that have not already been exploited in developing it. They correspond to the premises appearing to the left of a given turnstile in the sequent proofs discussed in the last section. Our aim in a developing a gap will thus always be to see whether the goal of the gap is entailed by its active resources. And this means that the situation depicted in Figure 2.2.2-1, which is explicit at the beginning of the derivation, will be replicated, although less explicitly, throughout the development of a derivation.

2.2.3. Rules for derivations

One way of developing a gap is to restate our problem so that one of its resources can be dropped from consideration, perhaps adding others of equivalent power but simpler form. We will call this process *exploitation*, and one example is provided by the way we implement the law for conjunction as a premise. That principle tells us that anything we can conclude from premises that include a conjunction can still be concluded if we replace the conjunction by its two components. In derivations, we will apply this idea by adding, as further resources, both of the conclusions that can reached from a conjunction by Ext. By adding both conclusions, we eliminate any further need to consider the conjunction we are exploiting; but, since both conclusions may not be needed to reach our ultimate goal, a derivation may contain conclusions that are never used later.

The derivation rule Extraction thus takes the form shown in Figure 2.2.3-1.

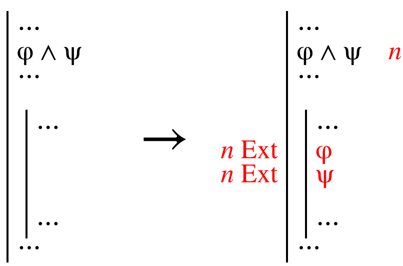

Fig. 2.2.3-1. Developing a derivation by exploiting a conjunction at stage *n*.

On the left, the gap is shown nested inside scope lines (two are shown but there may be just one or more than two). A conjunction is displayed at the top to show that it is among resources available for use in this gap. It is shown to the right of one of the scope lines running to the left of the gap but not to the right of the other. The requirement this illustrates is that a resource being exploited need not be inside all the scope lines to the left of the gap but cannot be inside any extra ones; that is, *all lines to the left of the resource being exploited must continue to the left of the gap it is exploited in*.

The right side of the figure illustrates the results of exploiting the conjunction. When we exploit it, we add its components as new resources at the top of the gap. If either component of the conjunction should happen to be already among the active resources of the gap, it would not be necessary to add this component again; but there is nothing wrong with doing so, and examples in the text will generally add it. (Although this practice may make the derivation slightly less compact, it makes it possible to focus solely on the parts of the derivation that are immediately relevant to the rule—i.e., the ones displayed in the diagram above.)

The number *n* of the new stage in the development of the derivation is written to the right of the conjunction to show that it has been exploited at this stage, and the stage number is also shown, along with the label Ext, to the left of each of the two lines that are added. Once the conjunction has been exploited, it is no longer an active resource for this gap though it could be active in other gaps (we will see later how to tell). The numbers in a derivation thus record the order of the development and also provide a way of telling when and where resources are exploited. These numbers are also one of the devices derivations use to encode the structure of tree-form proofs: they mark the relation between premises and conclusion that tree-form proofs marked the horizontal lines between premises and conclusions. In English argumentation, words and phrases like therefore, hence, and it follows that indicate the same sorts of connections in a less explicit way.

Exploiting resources like this is one way to narrow a gap. Another way to narrow a gap is to restate the problem it represents so that the goal we seek to reach is replaced by one or more simpler goals. We will call this process *goal planning*. The law for conjunction as a conclusion tells us how we may plan for a goal that is a conjunction. Such a goal is entailed by our active resources if and only if each of its components is entailed. So the project of reaching a conjunction $\varphi \wedge \psi$ from given resources comes to the same thing as completing two projects—namely, reaching each of the components φ and ψ from those same resources. This sort of goal planning thus uses Cnj and takes the form shown in Figure 2.2.3-2.

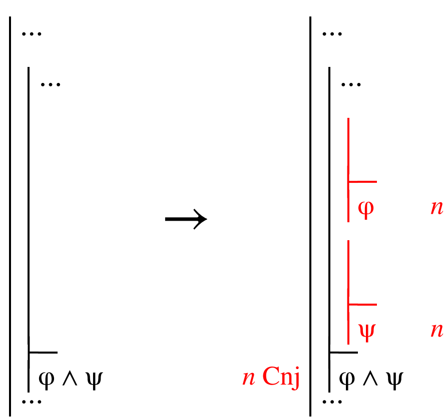

Fig. 2.2.3-2. Developing a derivation by planning for a conjunction at stage *n*.

On the left, no assumptions are made about the resources, but the goal is shown as a conjunction. On the right, we have introduced two new gaps, each with one of the conjunction's components as its goal. The two new goals bring with them two scope lines and are marked off by horizontal lines (as was the initial conclusion) to show that they represent the new material that led to the use of new scope lines. At the right of each of the new goals is a number showing the stage at which it was added. The same number appears to the left of the goal along with the label Cnj.

While in the case of Ext, numbers appeared at the left of the resources that were added and at the right of the resource being exploited, numbers here appear on the right of the new goals and at the left of the old one. This is because the new goals added by Cnj are introduced as premises from which the old goal may be concluded while the resources added by Ext are added as conclusions drawn from the resource that is exploited. Still, in both cases the numbers mark a connection between premises and conclusions. The numbers also show for both rules how an element of the derivation has been superceded by new additions. But, in the case of Cnj, this information is also provided by the added gaps: a gap will always have exactly one goal, and that goal will appear immediately below it.

The new gaps introduced in planning for a conjunction initially have the same active resources as the original gap. As resources are exploited in narrowing one of the gaps, these resources will become inactive for that gap; but they will remain active for the other gap until they are exploited there. When a derivation contains more than one gap, the question of where resources are active becomes important, and something will be said about it before too long. But, when we are dealing with conjunction alone, it is possible to exploit the initial resources completely before we plan for goals. As a result, a general discussion of active and inactive resources can be postponed until we have considered an actual example of a derivation.

What we cannot postpone is an account of how a gap may be closed. If the goal of a gap appears also among its resources, the law for premises tells us that the goal is entailed by these resources. That means we have succeeded in making a connection between our resources and that goal, and the gap may be closed. The rule we use to do this is shown in Figure 2.2.3-3 below.

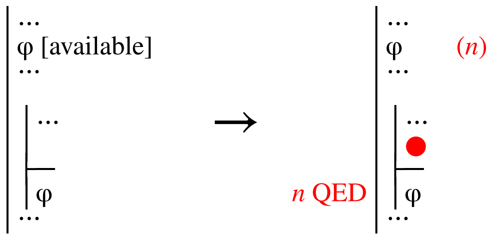

Fig. 2.2.3-3. Closing a gap by locating its goal among its resources.

The label for this rule abbreviates the Latin quod erat demonstrandum (which might be translated as what was to be proven), a phrase that is traditionally used when a planned conclusion is reached.

The stage number appears to the left of the goal (along with the label) since the goal is the conclusion, and it appears to the right of the resource since the resource is the premise. The latter number is enclosed in parentheses to indicate that the premise is not here being exploited. Since the gap is closed, the question whether a resource is active or not becomes moot; but this sort of notation will be used later in other cases where resources are used without being replaced by simpler resources of equivalent content, and QED shares with these rules the feature that the resources to which it is applied do not need to be active. To make it easy to see that the gap is now closed, we put the symbol ● (a *filled circle*) in it. This is really not part of the derivation itself and is not given a stage number; it instead functions like stage numbers to indicate the organization of a derivation. Think of an analogy with written language: the symbol ● marks the end of a series of stages in the way a period marks the end of a series of words.

2.2.4. An example

Now, let us look at an example using these rules. The development is shown stage by stage below. At each stage, new material is shown in red. Resources that are exploited or goals that are planned for are shown in blue. At each of the stages 1 and 2, a resource is exploited. The added resources are conclusions drawn from the exploited resource, so the number of the stage is written at the left of the resources that are added and at the right of the one that is exploited.

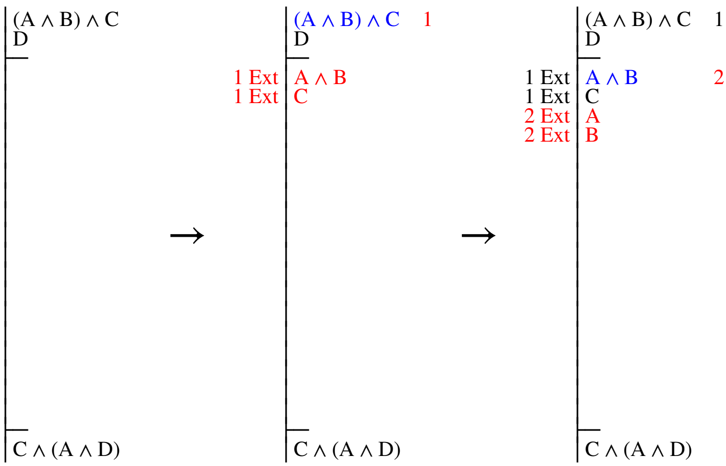

In stages 3 and 4, we plan for goals. The goals we add in each case are premises from which we plan to conclude the goal we are planning for. The stage number therefore appears at the right of the new goals and to the left of the old one.

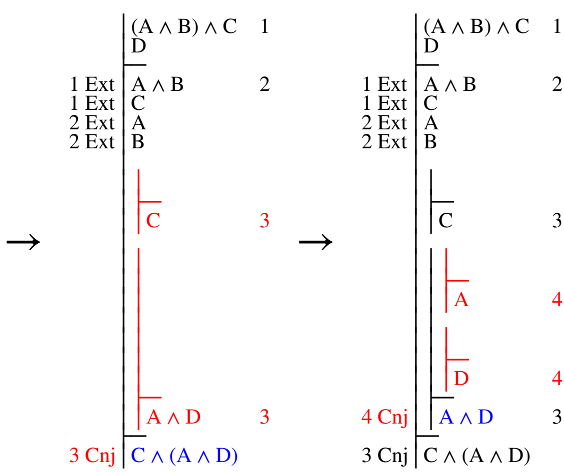

In the last three stages we close gaps. Although these are separate stages, they are independent of one another and could have been done in any order, so all three are shown together. No sentences are added and the stage numbers merely mark the connection between resources that serve as premises and the goals that are concluded from them (both shown in blue).

	$(A \wedge B) \wedge C$	1
	D	(7)
1 Ext	$A \wedge B$	2
1 Ext	C	(5)
2 Ext	A	(6)
2 Ext	B	
	●	
5 QED	C	3
	●	
6 QED	A	4
	●	
7 QED	D	4
4 Cnj	$A \wedge D$	3
3 Cnj	$C \wedge (A \wedge D)$	

If your browser has JavaScript enabled, the diagram below can be used to display each stage in the development of the derivation we have been considering.

	$(A \wedge B) \wedge C$	1
	D	(7)
1 Ext	$A \wedge B$	2
1 Ext	C	(5)
2 Ext	A	(6)
2 Ext	B	
	●	
5 QED	C	3
	●	
6 QED	A	4
	●	
7 QED	D	4
4 Cnj	$A \wedge D$	3
3 Cnj	$C \wedge (A \wedge D)$	

When this sort of animation is not available, the stage numbers in a completed derivation can be used to reconstruct its history.

Glen Helman 05 Aug 2010

2.2.5. Two perspectives on derivations

The locations of the stage numbers appearing in a derivation reflect the patterns of argument on which the derivation rules are based. The label for a rule always appears to the left of the conclusion of such an argument, and the number of the stage at which the rule was applied appears not only next to the label but also to the right of the premises of the argument. A tree form proof can be reconstructed from the derivation by beginning with the final conclusion and working backward to the premises from which it was concluded, the premises from which those were concluded, and so on.

If we apply this idea to the example of the last section (which is reproduced below), we get the tree-form proof following it.

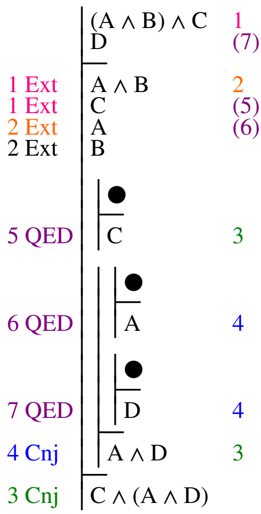

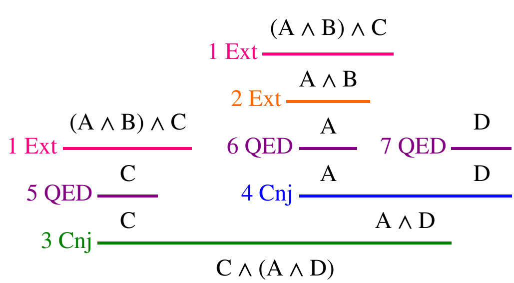

The sentence B concluded by Ext at the second stage of the derivation does not appear in the tree-form proof because it is not used as a premise for any later conclusions, something that can directly determined from the derivation by the fact that it has no stage number to its right.

Looked at in this way, a derivation could be thought of as the result of disassembling a tree-form proof and stacking the pieces up vertically. When reassembling the tree, we paid no attention to the horizontal organization provided by scope lines. The order of the stage numbers played no role either: they could just as well have been arbitrary codes used to mark corresponding parts of the tree so they could be fit together again. Indeed, even the vertical order of the lines of the derivation did not matter. Matching numbers on the left with numbers on the right is all that was necessary to reassemble the tree, and pieces could have been given to us in an unorganized heap. However, all these features of derivations, which are not needed to reconstruct a tree-form proof, do matter for another, and more important, way of looking at derivations, one in which a derivation is associated with a sequent proof.

To see this association, first use the stage numbers, scope lines, and the vertical ordering of lines to determine the way the gaps of the derivation develop over time, beginning with the intial gap, eventually dividing, and finally closing. That is shown on the right in the diagram below, where the stages are arrayed left to right and gaps are indicated by circles, with a filled circle used to indicate closure and an *empty circle* used to indicate a gap that is open. Colors are used to emphasize where and when changes occur.

	$(A \wedge B) \wedge C$	1
	D	(7)
1 Ext	$A \wedge B$	2
1 Ext	C	(5)
2 Ext	A	(6)
2 Ext	B	
	●	
5 QED	C	3
	●	
6 QED	A	4
	●	
7 QED	D	4
4 Cnj	$A \wedge D$	3
3 Cnj	$C \wedge (A \wedge D)$	

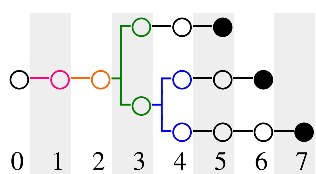

Now turn the tree pattern counterclockwise so that the tree grows upward. The result is shown on the left below. Then associate with each open gap the argument whose conclusion is the goal of the gap and whose premises are the active resources of the gap, a argument that we will refer to as the *proximate argument* of the gap. These two steps yield the tree on the right below (where the fact that an argument is valid is again indicated by a filled circle). Colored sentences are new additions as the tree grows (something that is shown by the dashed lines below them); the other sentences at a stage are repeated from earlier stages. Resources that become inactive and goals that are replaced as the derivation develops have dashed lines above them.

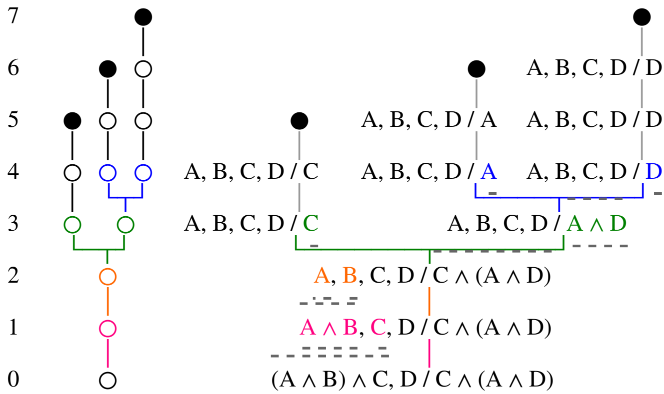

The tree on the right amounts to a schematic way of writing a sequent proof. The only differences are the use of the argument slash instead of the entailment sign and the more graphic indication of the branches. We will call this the *argument tree* associated with a derivation. It serves not only to emphasize the features derivations share with sequent proofs but also to present the sort of information about derivations that will be needed when we go on (in 2.3) to consider the general properties of derivations. Indeed, a derivation can be thought of as an abbreviated way of writing its argument tree.

2.2.6. More rules

A couple of the principles for $\top$ and $\bot$—in particular, with the laws for $\top$ as a conclusion and $\bot$ as a premise—have a role to play in derivations. Like the laws for conjunction, these laws have associated patterns of valid argument:

$$\text{ENV}\ \frac{}{\top} \qquad \text{EFQ}\ \frac{\bot}{\varphi}$$

The label for the second, EFQ, abbreviates the Latin ex falso quodlibet (which might be translated as from the false, whatever), a traditional way of stating the law for $\bot$ as a premise, and the label for the first, ENV, abbreviates ex nihilo verum (from nothing, the true), which gives a corresponding statement of the law for $\top$ as a conclusion.

The two other laws for $\top$ and $\bot$ do not have associated patterns of argument and will not be associated with steps in proofs. The law for $\top$ as a premise does not point to a pattern of argument whose conclusion could replace $\top$, for it tells us that $\top$ may simply be dropped from the premises. In fact, it will be just as easy to retain $\top$ as an active resource but ignore it. And that will make our handling of $\top$ more like our handling of $\bot$. For we cannot apply the principle for $\bot$ as an alternative unless we begin with multiple alternatives or end with none, so it is not a principle of entailment at all and provides no way of replacing $\bot$ as a conclusion. This does not mean that $\bot$ as a conclusion is insignificant in the way $\top$ is insignificant as a premise, but the role of $\bot$ as conclusion is to mark an entailment as a claim of inconsistency, and such claims will be established by applying principles to their premises rather than to their conclusion. (However, we will eventually have some rules for exploiting resources that we will apply only when the goal is $\bot$.)

The principles ENV and EFQ figure in derivations as rules for closing gaps. In the case of the first, it is enough for a gap to be closed that it have $\top$ as its goal. No resource is involved, and the stage number appears only as an annotation to the goal.

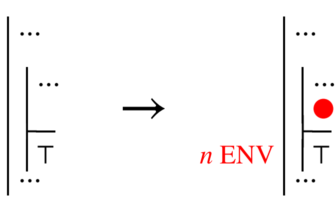

Fig. 2.2.6-1. Closing a gap that has $\top$ as its goal.

The rule EFQ takes a form much like QED.

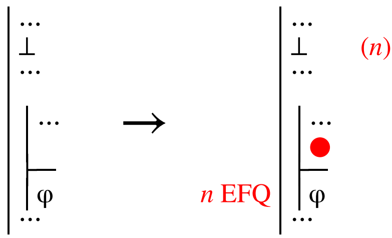

Fig. 2.2.6-2. Closing a gap that has ⊥ among its resources.

The difference is that having ⊥ as a resource enables us to close a gap no matter what its goal is. (If the goal also was ⊥, either EFQ or QED could be used.)

Here are examples of the use of these rules:

	A	(5)
	B	(3)
	●	
3 QED	B	2
	●	
4 ENV	⊤	2
2 Cnj	B ∧ ⊤	1
	●	
5 QED	A	1
1 Cnj	(B ∧ ⊤) ∧ A	

	A ∧ (⊥ ∧ B)	2
2 Ext	A	
2 Ext	⊥ ∧ B	3
3 Ext	⊥	(4),(5)
3 Ext	B	
	●	
4 EFQ	C	1
	●	
5 EFQ	D	1
1 Cnj	C ∧ D	

Notice that, while every stage number of the second derivation appears somewhere among the annotations on its right-hand side, the same is not true of the first derivation because stage 4 is missing. Of course, that's because stage 4 is when we used ENV, and ENV is a valid argument without premises. Since stage numbers appearing in annotations on the right-hand side of a derivation mark the use of a line as a premise and ENV is the only form of argument we have seen so far that has no premises, a use of the rule ENV should be the only reason for a stage number to appear on the left of a derivation but not on the right.

You can use this idea as a way of checking for errors, and there are some further generalizations like this that you can use as checks. We will have no rules without conclusions, so every stage number should appear somewhere in the left-hand annotations. And, in a completed derivation whose gaps all close, all sentences other than assumptions (which, for now, are just the initial premises) will be conclusions and thus should have annotations on their left-hand side. Resources that are never used may appear with no annotations on their right; but, as you are constructing a derivation, it can be very useful to check for the absence of right-hand annotations because this can lead you to notice resources that you have not yet exploited. And, when we go on (in 2.3) to use derivations to show that claims of entailment fail, a check for the absence of right-hand annotations will be the key test of whether we done everything possible to complete a derivation.

2.2.7. Resources

The ideas of available and active resources have been used at several points already, but they have not yet been explained fully. A resource counts as *available* in a gap if it was entered either as one of the initial premises of the derivation or in the course of developing the gap in question. The system of scope lines can be used to tell which resources are available in a gap: a resource is available if every scope line to its left continues unbroken at the left of the gap.

One way of thinking about this is shown in Figure 2.2.7-1.

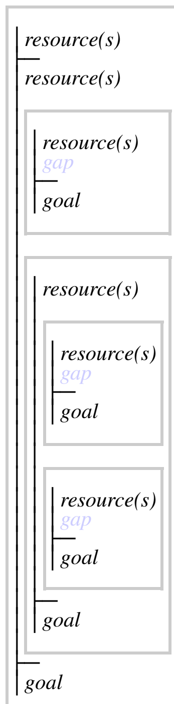

Fig. 2.2.7-1. The boxes indicated by the scope lines of a derivation. If JavaScript is enabled on the browser you are using, moving the cursor over a resource will color the gaps in which it is available green and shade areas where it is unavailable. Moving the cursor over a gap will color resources available in it green and shade areas whose resources are unavailable to it. The resource or gap that the cursor is over will be colored blue and underlined.

You may suppose that each scope line indicates the left side of a box and that a resource is available only to the gaps that are also within the smallest box containing it.

A resource is *active* in a gap if it is available in that gap and has not already been exploited in narrowing it. The easiest way to locate the active resources of a gap is to scan the available resources and eliminate the inactive ones. To be inactive in any gap, a resource must have been exploited at some stage. If it has, there will be an unparenthesized stage number to its right. A resource may have been exploited only in some gaps and may still remain active in others. To be inactive in a given gap, the resource must have been exploited in narrowing the gap. To see whether this is so, we need to check all resources and goals that were introduced at a stage when the resource was exploited (i.e., at a stage whose number appear unparenthesized to the resource's right). (So far, we have seen goals introduced only in the course of planning for more distant goals, but in later chapters they will be introduced as part of the exploitation of certain resources.) If any such resource or goal is such that the smallest box containing it also contains the gap we are considering, it was introduced in the course of developing the gap. A resource may be exploited more than once, so there may be several stage numbers you will need to check. If any of them was a stage in which the gap you are considering was developed, the resource is no longer among the active resources of the gap. This description of the process may make it sound rather daunting, but in practice you will find that it is usually obvious which resources have been exploited in developing a given gap.

The partially developed derivation shown below has been designed to provide an example of a resource that has been exploited without being exploited in all gaps in which it is available.

	(A ∧ B) ∧ C	1
1 Ext	A ∧ B	3
1 Ext	C	
3 Ext	A	
3 Ext	B	
	A	2
	B	4
	C	4
4 Cnj	B ∧ C	2
2 Cnj	A ∧ (B ∧ C)	

The three steps at the top of the derivation are resources available for each of the derivation's three gaps. The first, (A ∧ B) ∧ C, is inactive in all three gaps. It was exploited at stage 1, and that was the initial stage of development for all the gaps of the derivation. The second resource, A ∧ B, is inactive for the first of the gaps (having been exploited at stage 3 in developing this gap), but it is active for the remaining two gaps since the resources introduced at stage 3 did nothing to narrow these gaps (as is shown by the fact that the gaps are outside the smallest box surrounding the resources with 3 at their left). The third resource C has not been exploited at all (and could not be since it is not a conjunction), so it is active for all three gaps. Since the resource exploited at stage 3 must be exploited again in order to close the second gap, it would have been a little more efficient to exploit this resource before dividing the initial gap in two; but the derivation as shown is perfectly correct (though still unfinished).

You may suppose that a given gap can see only those parts of a derivation that are not boxed off from it—i.e., only those parts all of whose scope lines continue to the left of the gap. If a stage number appears at the left only in parts of the derivation that are invisible to the gap, this stage number is also invisible—even when it appears to the right of resources that are visible.

This idea is illustrated in Figure 2.2.7-2 below where the same derivation is shown from the perspective of each of the three gaps in turn.

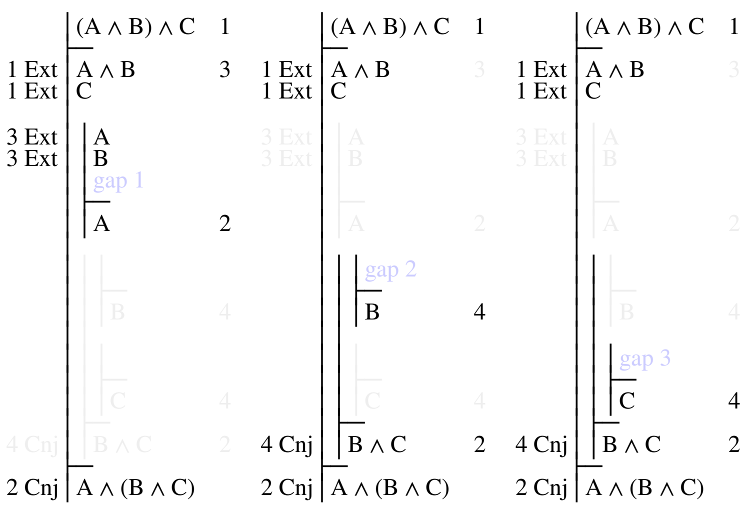

A B C

Fig. 2.2.7-2. A derivation from the perspective of each of its three gaps.

Material that is boxed off from a gap is shown in very light gray. Notice that the number 3 at the right of the second line is invisible to the second and third gaps. As we saw earlier, that is because all the development at stage 3 is boxed off from the second and third gaps.

Any derivation can be thought of as the result of superimposing separate layers like these. There will be one layer for each gap with a gap's layer depicting its perspective on the derivation. This corresponds directly to a feature of argument trees: a gap can see what is on the path from it back to the root of the tree, and the superimposing layers to make up a derivation corresponds to superimposing paths to make up a tree. When we distinguish the resources available for a gap or determine whether a resource has been used to narrow a gap, we are really considering that gap's layer separately, which is to say we are considering its path to the root apart from paths that have branched off.

When a gap is divided before a resource is exploited to narrow it, it is possible to exploit the resource to narrow several gaps at once. This is shown in the partial derivation below (which has the same initial premises and conclusion as the one we have been considering).

```
       | (A ∧ B) ∧ C     1
       |----
1 Ext  | A ∧ B           4
1 Ext  | C
       |
4 Ext  | | A
4 Ext  | | B
       | |
       | |----
       | | A             2
       |
4 Ext  | | | A
4 Ext  | | | B
       | | |
       | | |----
       | | | B           3
       | |
       | | |
       | | |----
       | | | C           3
       | |----
3 Cnj  | | B ∧ C         2
       |----
2 Cnj  | A ∧ (B ∧ C)
```

In this derivation, one of the resources has just been exploited at stage 4 to narrow two different gaps. Thereafter, it is inactive in these gaps but still active in the third (where it happens to be unneeded). Some of the resources added at stage 4 will be invisible to each of the first two gaps; but, because other added resources are visible, the number 4 at the right is visible from both these gaps. However, none of the resources added at stage 4 is visible from the third gap, so the number 4 at the right is not visible from it.

Since we use a similar numerical notation for both resources that are exploited and goals that have been planned for, you might expect that the concepts of availability and activity can be applied to goals as well as resources; and, indeed, they can be. If we were to consider derivations for conditional exhaustiveness, we would need to engage in the same sort of accounting for goals that we have been considering for resources. However, in a system of derivations for entailment alone like the one we will actually use, each gap has just one active goal, which appears just below the gap. Goals at earlier stages of a gap's development (i.e., the goals that are not boxed off from the gap) could be described as "available", but they are not available for any sort of use. In particular, although we can consider all available resources when looking for a way of closing a gap, it is only the active goal and not any earlier one that we consider. (Some of the arguments of 2.3.3 could be used to show that considering all "available" goals would not lead us to count an invalid argument as valid, but looking at derivations in this way would make them less like the patterns of ordinary explicit deductive argumentation, which seem to be focused always on a single conclusion.)

2.2.s. Summary

1 Principles of entailment can be applied in concert by using the graphical idea of a tree—and in more than one way. Tree-form proofs provide a natural notation for applying the valid patterns of argument Cnj and Ext. Alternatively, we can consider principles of entailment that state conditions for the validity of arguments that have conjunctions as conclusions or as premises. These, too, can be combined in a tree as a sequent proof to show an argument is valid. A sequent proof can also be thought of as a tree whose growth traces the conditions that must hold for the argument at its root to be valid.

2 In fact, we will use a different, more compact notation for combining principles of entailment—a kind of natural deduction system that we will refer as a system of derivations. This notation presents the project of showing that an entailment holds as the task of closing a gap between its conclusion, which serves as a goal, and its premises, which serve as resources. As we narrow the initial gap (and others that result from it), we develop the derivation. Derivations also have a tree structure displayed in a system of vertical scope lines which indicate the resources and goals relevant to various parts of the derivation.

3 The laws of entailment appear as rules for exploiting resources, planning for goals, and closing gaps. There are rules for each of the patterns of argument that figure in tree-form proofs. The key rules for conjunction are Extraction (Ext) and Conjunction (Cnj). Quod Erat Demonstrandum (QED) is used to close a gap when its goal is among its resources, and the symbol ● (a filled circle) marks a closed gap.

4 When a derivation is developed, numbers are used along with the labels for rules to record both the order of the development and the connection between the premises and conclusions of the rules.

5 The branching structure of tree-form proofs is replicated in derivations by the system of cross-references provided by stage numbers. And the branching structure of sequent proofs lies in the way gaps develop, something indicated by the order of stage numbers and the arrangement of scope lines. This structure, together with the proximate argument of each gap (formed from its active resources and its goal), forms an argument tree.

6 Principles of entailment for other logical forms will be associated with further rules. Those for ⊤ and ⊥ are the rules Ex Nihilo Verum (ENV) and

Ex Falso Quodlibet (EFQ), which figure in derivations as rules for closing gaps.

7 We keep track of changes in the information contained in goals and resources by using the scope lines of a derivation to tell in which gaps given resources are available and in which gaps available resources are still active.

2.2.x. Exercise questions

1. Restate the derivation below in two ways: (i) as a tree-form proof, labeling each horizontal line with the number of the stage at which it is entered, and (ii) as its associated argument tree. That is, do with it what is done with the example in 2.2.5 (ignoring the extra decoration, such as colors and dashed lines, that appeared there).

	$(A \wedge C) \wedge B$	1
1 Ext	$A \wedge C$	2
1 Ext	B	(4)
2 Ext	A	
2 Ext	C	(5)
	●	
4 QED	B	3
	●	
5 QED	C	3
3 Cnj	$B \wedge C$	

2. Use the system of derivations to establish each of the following claims of entailment:

a. $A \wedge B \vDash B \wedge A$

b. $A \vDash A \wedge A$

c. $A \wedge (B \wedge C) \vDash (C \wedge B) \wedge A$

d. $A, B \wedge C, D \vDash (C \wedge (B \wedge A)) \wedge B$
[The derivation for **d** will have three premises above the initial horizontal line.]

e. $A \wedge (B \wedge C) \vDash (B \wedge A) \wedge (C \wedge A)$

For more exercises, use the exercise machine.

Glen Helman 03 Aug 2010

2.2.xa. Exercise answers

1.

$$\text{3 Cnj}\;\dfrac{\text{4 QED}\;\dfrac{\text{1 Ext}\;\dfrac{(A \wedge C) \wedge B}{B}}{B} \qquad \text{5 QED}\;\dfrac{\text{2 Ext}\;\dfrac{\text{1 Ext}\;\dfrac{(A \wedge C) \wedge B}{A \wedge C}}{C}}{C}}{B \wedge C}$$

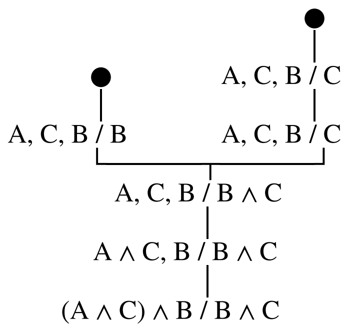

2. **a.**

	$A \wedge B$	1
1 Ext	A	(4)
1 Ext	B	(3)
	●	
3 QED	B	2
	●	
4 QED	A	2
2 Cnj	$B \wedge A$	

b.

	A	(2),(3)
	●	
2 QED	A	1
	●	
3 QED	A	1
1 Cnj	$A \wedge A$	

c.

	$A \wedge (B \wedge C)$	1
1 Ext	A	(7)
1 Ext	$B \wedge C$	2
2 Ext	B	(6)
2 Ext	C	(5)
	●	
5 QED	C	4
	●	
6 QED	B	4
4 Cnj	$C \wedge B$	3
	●	
7 QED	A	3
3 Cnj	$(C \wedge B) \wedge A$	

d.

	A	(7)
	$B \wedge C$	1
	D	
1	B	(6)
1	C	(5)
	●	
5 QED	C	3
	●	
6 QED	B	4
	●	
7 QED	A	4
4 Cnj	$B \wedge A$	3
3 Cnj	$C \wedge (B \wedge A)$	2
	●	
	B	2
2 Cnj	$(C \wedge (B \wedge A)) \wedge B$	2

e.

	$A \wedge (B \wedge C)$	1
1 Ext	A	(7),(9)
1 Ext	$B \wedge C$	2
2 Ext	B	(6)
2 Ext	C	(8)
	●	
6 QED	B	4
	●	
7 QED	A	4
4 Cnj	$B \wedge A$	3
	●	
8 QED	C	5
	●	
9 QED	A	5
5 Cnj	$C \wedge A$	3
3 Cnj	$(B \wedge A) \wedge (C \wedge A)$	

Glen Helman 03 Aug 2010

2.3. Failed proofs and counterexamples

2.3.0. Overview

Derivations can also be used to tell when a claim of entailment does not follow from the principles for conjunction.

2.3.1. When enough is enough
A derivation is stopped only when no more rules can be applied. When that is so, any open gap has reached a dead end.

2.3.2. Dividing gaps
The active resources of any dead-end gap can be divided from its goal. To put it another way, we have enough rules to develop further any gap whose proximate argument cannot be divided.

2.3.3. Validity through the generations
If we describe as descendents of a gap the gaps that result from developing and perhaps branching it, the validity of the proximate argument of a gap rests on the validity of the proximate arguments of its descendents.

2.3.4. Sound and safe rules
The derivation rules are designed so that, if a gap can be divided, so can at least one descendent at every stage and, moreover, all of its ancestors.

2.3.5. Presenting counterexamples
Because we have enough rules and the ones we have are well-behaved, any gap that reaches a dead end shows us how to divide the premises of the initial argument from its conclusion.

2.3.6. Reaching decisions
A derivation will always reach a point where we must stop either because all gaps are closed or because there is an open gap to which no more rules can be applied.

2.3.7. Soundness and completeness
The properties of this system of derivations combine to show that it establishes the validity of no argument that is not valid and does establish the validity of all that are.

2.3.8. Formal validity
The sort of validity we test with derivations is the general validity of arguments with a given form. An argument that is not valid in virtue of a given form could be valid nonetheless, and its validity may be recognized by a deeper analysis of its form.

2.3.1. When enough is enough

So far we have seen only derivations whose gaps all close, derivations which show that arguments are valid. But not all arguments are valid, so there ought to be derivations whose gaps do not all close. If there is no point at which the gaps of a derivation all close, we will eventually have to give up work on it even though it still has open gaps. So we should ask what might lead us to give up work and what, if anything, we can conclude if we do have to stop.

The short answer to the first of these two questions is that we must give up on a derivation when we run out of rules to apply, either to develop a gap or close it. Here's a simple example of a derivation for which that has happened.

	$(A \wedge \top) \wedge B$	1
1 Ext	$A \wedge \top$	6
1 Ext	B	(4)
2 Ext	A	
2 Ext	$\top$	
	●	
4 QED	B	3
	○	$B, A, \top \nvDash C$
	C	3
3 Cnj	$B \wedge C$	

The gap that is marked with the empty circle ○ has C as its goal, and we currently have no rule to plan for such a goal. There are conjunctions among the available resources of the gap; but they were exploited in the course of developing this gap, so they are no longer active. Also, none of the rules for closing gaps apply here: not QED because the goal is not one of available resources, not EFQ because $\bot$ is not a resource, and not ENV because the goal is not $\top$. In short, no rule of any of the three sorts can be applied at this point. Notice that the resources added by exploiting $A \wedge \top$ at stage 2 were never used later (hence there are no line numbers to their right). As a result, this exploitation could have been postponed the end. However, the resource $A \wedge \top$ must be exploited before we end work on the derivation because, until it is exploited, there is a way of developing the derivation further.

We will describe an open gap to which no more rules apply as a *dead-end gap*. (Although the qualification dead-end will be reserved for open gaps—indeed, a gap that has been closed is in one sense no longer a gap—we will often speak somewhat redundantly of “dead-end open gaps.”) In these terms, we can say that we are forced to abandon a derivation when every open gap has reached a dead end. When we consider the significance of dead-end open gaps, we will see that we *may* abandon a derivation as soon as one open gap has reached a dead-end. As in the example above, we will use the empty circle to mark open gaps that have reached a dead end and are thus permanently open. And, also as is done in that example, to the right of this sign, we will use the sign ⊭ (*negated double right turnstile*) to say that, with respect to the analysis of them displayed in the derivation, the active resources do not entail the goal. (The reason for qualifying this by reference to the displayed analysis will be discussed in 2.3.8.)

The way the gaps have developed in this derivation is shown in the following tree:

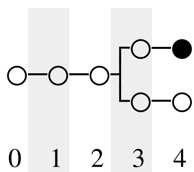

The gap that remains open at the end had reached a dead end at stage 3, but it is shown to continue at the next stage because it remains open as the derivation develops elsewhere. As we will see, a single dead-end gap in a derivation for a claim of entailment tells us that the claim fails, so work *may* be stopped as soon as a dead-end is reached. But there is nothing wrong with continuing as long as there are rules to be applied to other gaps, and we will often do so in examples. In general we will not assume that a derivation stops as soon as there is a dead-end gap, so to say that gap has reached a dead-end is not to say that it does not continue at later stages; it is to say rather that we can be sure it will never close.

From one point of view, the function of a derivation is to transform the question whether an argument is valid into an analogous question about one or more simpler arguments. This is the aspect of a derivation that is displayed in the growth of its argument tree, which is shown below for the argument we have been considering.

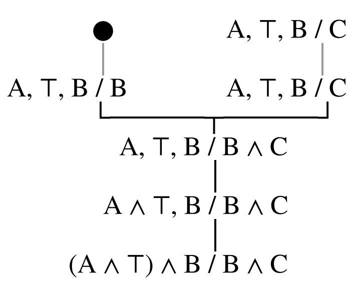

The proximate argument of a dead-end open gap is the end of the line in this process; it will not be developed further though it may be repeated. We will call the argument whose validity we initially asked about, the one at the root of the tree, the *ultimate argument* of the derivation. It is the proximate argument of the initial gap of the derivation. The contrast between the proximate argument of a gap and the ultimate argument of a derivation is the source of our use of the term proximate: the proximate argument of a gap is our *immediate* concern while our *final* goal is to decide whether the ultimate argument of the derivation is valid.

In discussing the significance of dead end gaps, we will look first at what reaching a dead-end tells us about the proximate argument of the gap that has stopped developing and then consider the connection between the validity of the ultimate argument of a derivation and the existence of dead-end gaps. In terms of the argument trees, this means we will look first at the tips of unclosed branches and then ask about the connection between the tips of branches and the root of the tree.

2.3.2. Dividing gaps

Now, let's look more closely at what we can say in general about the significance of dead-end open gaps. First of all, recall what led us to conclude that the gap in the example of the last section could not be developed further. A dead-end gap must not have a conjunction either as its goal or among its active resources, for otherwise we could apply the rules Cnj or Ext. Moreover, it must not have $\top$ as a goal or $\bot$ as a resource, or else we could apply the rules ENV or EFQ. Finally, its goal must not be among its resources because then we could apply the rule QED. So the active resources of dead-end gaps are limited to unanalyzed components and $\top$ and their goals are limited to unanalyzed components and $\bot$; and no dead-end gap can contain an unanalyzed component both as an active resource and as its goal.

This means that we can assign truth-values to the unanalyzed components appearing in a dead-end gap in a way that makes its active resources true and its goal false. Since no unanalyzed component appears both as a resource and as the goal, we can make any that appears as a resource **T** and any that appears as the goal **F**. While we are not free to assign values to $\top$ and $\bot$, the first can appear only as a resource and the second only as the goal so they will not keep us from having true resources and a false goal. In short, we can assign truth values in a way that divides the proximate argument of the dead-end gap.

In noting this, we described an assignment of truth values to unanalyzed sentences. This is an extensional interpretation in the sense discussed in 2.1.8, and it can be presented in a table. The following table displays the interpretation defined by the dead-end gap of the example we have been considering.

A	B	C	B,	A,	$\top$	/ C
T	T	F	Ⓣ	Ⓣ	Ⓣ	Ⓕ

The extensional interpretation of unanalyzed components appears on the left of the table. On the right are the resulting truth values of resources and goals of the gap (which mainly just repeat the assignments). No value is assigned to $\top$ on the left because its truth value is stipulated by the meaning of the sign. Unlike A, B, and C, the sentence $\top$ is not something whose value we are free to assign, and it is something that has a value without any assignment being made by us.

The idea of division that was introduced in 1.4.2 can be extended to speak in a compact way of what this interpretation does. When an interpretation divides the active resources of a gap from its goal—that is, when it divides the proximate argument of the gap—we will say that it *divides* the gap. If there is some interpretation that divides a gap, we will say the gap is *divisible*; otherwise we will say that it is *indivisible*. So an indivisible gap is one that has a valid proximate argument, and a divisible gap is one whose proximate argument is not valid. Note also that an extensional interpretation which divides a gap counts as a counterexample to the validity of the proximate argument of the gap (where the validity we speak of is again validity relative to a particular analysis of the argument).

Although we certainly have more to show before we know that the system of derivations does what it is supposed to, we can say already that it has enough rules in a certain sense, for we know that, whenever the proximate argument of a gap is valid, some rule can be applied to either develop or close the gap. For if there is no rule allowing us to develop the gap, it has reached a dead end, and we have just seen that the proximate argument of a dead-end gap is not valid. We will indicate this sort of completeness in our rules by saying that a system of derivations is *sufficient* when every dead-end open gap is divided by some extensional interpretation. Of course, in saying that system is sufficient, we do not say that every gap whose proximate argument is invalid has already reached a dead end. We would not expect this to be true since it would mean that we would never need to apply any rules at all in the case of an invalid argument. Indeed, one of the things we have yet to show is that any gap whose proximate argument is invalid will eventually reach a dead end.

2.3.3. Validity through the generations

The connection between the proximate arguments of dead-end gaps and the ultimate argument of a derivation lies in the properties of the rules for developing and closing gaps. We will begin to look at these properties in this section and then look at them more closely in the next.

It will help to have some ways of talking about the relations between gaps at various stages of a derivation. It is common to extend some genealogical vocabulary from family trees to trees in general. In our use of this vocabulary, we will say that any gap that results from applying a rule is a *child* of the gap to which the rule is applied and that the latter gap is its *parent*. It will be convenient to apply the same terminology to gaps that continue unchanged while others develop: a gap at one stage that is open but unchanged at the next stage is understood to have a single child. Looking farther up or down a line of descent, we will say that some gaps are *ancestors* or *descendants* of others. So in the tree of gaps associated with the derivation discussed in 2.2.5,

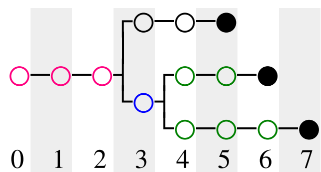

the lower gap at stage 3 has the gap at stage 2 as its parent and both that and the two earlier gaps as ancestors. Its children are the lower two gaps at stage 4 and its further descendants are the gaps to their right. The line of gaps at the top are neither ancestors or descendants of the gap in question.

In this terminology, the initial gap of a derivation is an ancestor of all gaps of all gaps at each later stage in its development; and they are all its descendants. Only open gaps will be part of these genealogies, so a gap that is closed at the next stage of its development has no children. Dead-end open gaps continue to have children if the derivation is continued at later stages (remember it need not be), yet they have reached a dead end in the sense that these children are always identical to their parents.

Next, let us develop a way of speaking about the effect of derivation rules on the distribution of valid and invalid arguments in the argument tree of a derivation. In the case of QED, we will initially limit ourselves to its use to close a gap whose goal is also among the active resources; the wider use of QED, to close gaps whose goals are among their available but inactive resources, will be considered in the next section.

The derivation rules Ext and Cnj are based on principles of entailment which give necessary and sufficient conditions for an entailment to hold. That is, each principle gives a list of conditions all of which must hold if a given entailment is to hold and which together are enough to insure that it holds. It may seem odd to say the same about the unconditional claims of entailment that lie behind the rules QED, ENV, and EFQ; but, by asserting an entailment unconditionally, they say that an empty list of conditions is sufficient for its truth (and, since an empty list cannot have a member that fails to hold, satisfying the list is trivially necessary since it is bound to be satisfied).

Phrased in terms of arguments, each principle tells us that a certain sort of argument is valid if and only each member of a (perhaps empty) list of arguments is valid. When the corresponding rule is applied to a gap, the gap is provided with children whose proximate arguments are those on the list (so the gap is given no children—that is, it is closed—if the list is empty).

rule	*prox. arg. of parent*	*prox. args. of children*
Cnj	$\Gamma / \varphi \wedge \psi$	Γ / φ Γ / ψ
Ext	$\Gamma, \varphi \wedge \psi / \chi$	$\Gamma, \varphi, \psi / \chi$
QED	$\Gamma, \varphi / \varphi$	(none)
ENV	$\Gamma / \top$	(none)
EFQ	$\Gamma, \bot / \varphi$	(none)

This means that the proximate argument of a gap to which a rule is applied is valid if and only if all the proximate arguments of any children it has are valid. And, of course, the same is true of a parent which acquires a child when the derviation is developed elsewhere because then there is only one child and its proximate argument is the same as its parent's.

To say that the proximate argument of a gap is valid is to say that the gap is indivisible, so we can say that a gap before the last stage is indivisible if and only if each one of any children it has is indivisible. It is usually more convenient to speak of divisibility (i.e., of the invalidity of the proximate argument), and we can rephrase what we have been saying in these terms as follows.

> A gap followed by another stage is divisible if and only if it has a child that is divisible.

This gives us necessary and sufficient conditions for the divisibility of a gap in terms of divisibility at the next stage, but it is stated only for cases where there is a following stage (though it does not require that the gap have children) and it is stated only for the immediately following stage. We will go on to consider what can be said of any gap and said with respect to any following stage. That will be enough to tie the divisibility of the initial gap with the state of the derivation after all work is done.

First note what we can say in cases where there are two stages following a gap. For a gap to be divisible in such circumstances, it must have a divisible child, which must itself have a divisible child. That is, a necessary condition for divisibility when there are two following stages is having a divisible grandchild. And that is clearly also sufficient, for a divisible grandchild will have a divisible parent, which will be a divisible child of the grandparent gap. Of course, the same thing will work for great-grandchildren, great-great-grandchildern, and so on, provided there are enough following stages.

In general, we can say this:

> For any pair of stages, one earlier than the other, a gap at the earlier stage is divisible if and only if it has a divisible descendant at the later stage.

Notice that this not only ties the divisibility of a gap to the divisibility of its descendants, however distant, but also holds for a gap when there are no later stages at all. The latter point is analogous to one made above about gap-closing rules: a generalization about an empty collection is bound to be true, no matter what it says, because there is nothing to serve as a counterexample.

These points are illustrated in the diagram below. It shows a sort of schematic argument tree that does not display actual arguments, only their validity or invalidity—i.e., their indivisibility or divisibility. It is intended to depict a derivation that has come to an end, so the one gap that remains open at the top is a dead end.

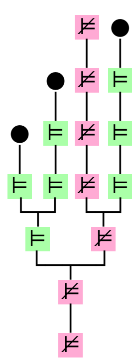

We can distinguish three sorts of cases in this tree. First of all, we know from the last section that the dead-end gap is divisible. It has no divisible descendent, but it is not a counterexample to the generalization above because there is no later stage. Next, all ancestors of the dead-end gap, right down to the root of the tree, are divisible because each has a divisible descendant. And finally, in the case of any of the other gaps—i.e., the ones whose proximate arguments are valid—there is a following stage (the last stage of the derivation if not an earlier one) at which the gap has no descendant at all, and so certainly has no divisible descendant. Also, notice that, at stages where such a gap does have descendants, all its descendents are indivisible.

There is a fourth sort of case that does not appear here, a gap that has no descendants but has not been closed and is not at a dead end. But this case will appear only in the last stage of an incomplete derivation, and the generalization says nothing about it because there is no later stage.

The generalization we have been considering tells us that the way we have taken the results of a derivation is correct. If there is a dead-end gap—and thus, by sufficiency, a divisible gap—the initial gap must be divisible, so the ultimate argument is invalid. On the other hand, if all gaps close, there is a stage (the one at which the last gap closes) at which the initial gap has no descendants, so it must be indivisible and the ultimate argument must be valid. Although this generalization does represent an important property of the system of derivations, we will not label it (in the way we have labeled the property of sufficiency) because we will go on in the next section to look further at the basis for this property and state (and label) some related properties that can be applied to a wider range of rules, including the extended use of QED that we excluded from consideration here.

2.3.4. Sound and safe rules

The necessary and sufficient conditions for divisibility and indivisibility developed in the last section were based on connections between the divisibility of gaps at successive stages. In this section, we will look more closely at the rules and consider not merely how the existence and non-existence of dividing interpretations is preserved as we develop a derivation but indeed how any dividing interpretations are themselves preserved. This closer look at the effect of rules will enable us to give an account of a wider range of possible rules, including the extended use of QED that was not covered in our discussion in the last section.

We begin by considering two properties a rule R might have:

R is *strict*	*when*	any interpretation of the derivation that divides a gap to which the rule R is applied also divides some child of the gap
R is *safe*	*when*	any interpretation of the derivation that divides a child of a gap to which the rule R is applied also divides the parent gap

When a rule is strict we never lose any gap-dividing interpretations as we apply the rule. When it is safe, we never gain any interpretations. It is the safety of our rules that implied that the condition for divisibility discussed in the last was sufficient while their strictness is the source of its necessity. In both cases, we generalize about interpretations of the whole derivation because an interpretation that divides a child gap need not assign truth values to enough sentences to count as an interpretation of the parent. However, every way of the interpreting the vocabulary of the proximate argument of a gap can be found in some interpretation of the derivation as a whole, so the restriction to interpretations of the whole derivation does not really limit the scope of the generalizations.

Although their association with the necessity and sufficiency of the same condition suggests a kind of parallel between them, these two properties do not have the same importance. Although we will see that strictness is a little more than we need to ask, any serious departure from strictness would undermine the central function of proofs: to establish validity. For then all gaps of a derivation might close even though the original argument was invalid. An unsafe rule would analogously undermine the use of derivations to establish invalidity because it would introduce the possibility that a derivation for a valid argument could lead us to a dead-end. But the role of derivations in establishing invalidity is less central, and their full use in that way depends also on a property (discussed in 2.3.7) that will fail for rules to be considered in the last two chapters. This means that safety is dispensible, but no viable system of proof could completely dispense with strictness.

Moreover, moves corresponding to unsafe rules are an important part of explicit deductive reasoning. For example, a natural approach when we seek a way to prove a mathematical result is to introduce a lemma (in the sense is discussed in 1.4.6) as a stepping stone to a final result. If the lemma represents a significant step beyond the premises, it may be no more obviously a valid conclusion from the premises than is the final conclusion we hope to establish. The introduction of such a lemma can be described as a conjecture, and this conjecture may be wrong: the lemma may not be a valid conclusion from our premises even when the final conclusion is valid. In short, by seeking to reach our conclusion by way of this lemma, we may be entering a blind alley. This is just the sort of thing that would appear in the context of derivations as a dead-end open gap in a derivation whose initial argument is valid. So conjecturing a lemma can be thought of as a step in discovering a proof that is valuable but unsafe.

Another step in a proof that can be valuable but is unsafe is a decision to focus on only some of the information in one's premises. This might seem quite different from a conjecture; but, combined with rules we will consider in the next chapter, a rule allowing us to conjecture a conclusion could lead us into a situation in which the active resources entailed less than did the resources at an earlier stage with the same goal.

Our interest in deductive reasoning is somewhat different from a mathematician's. We are aiming not at new and surprising conclusions but instead at fuller understanding of the steps by which deductive conclusions are reached. Consequently, we will not be considering the large deductive steps for which conjecturing lemmas is the only practical approach. We will make use of lemmas—and we will look at rules for doing so in 2.4—but the chief value of lemmas for us lies in a restricted range of cases where we can be sure that they are safe.

Earlier, we set aside uses of QED in which the goal of the gap we close is among its available resources but not among the active ones. To discuss such uses of QED, we need to consider a requirement that is less unyielding than strictness. The following property of a rule *R* is the one we will employ:

R is sound	*when*	any interpretation that divides both a gap to which the rule *R* is applied and all ancestors of this gap also divides some child of the gap

The difference lies in the added phrase and all ancestors of this gap. The addition makes soundness apparently weaker than strictness because, for soundness, we do not require that an interpretation divide a child gap simply because it divides the parent but only when it also divides all ancestors of the parent. However, when all rules are safe, a rule that is sound is also strict. For, when all rules are safe, an interpretation that divides a gap will also divide all ancestors of the gap. Thus, when there is a difference between soundness and strictness, it lies in their handling of the spurious dividing interpretations introduced by unsafe rules: with strict rule, such interpretations will continue to divide descendants while, with a sound rule, they might not. So a strict rule would force us to bear the burden of proving an unsafe conjecture while a sound rule might allow us to substitute a different way of reaching our initial goal.

And even when not all rules are safe, soundness is enough to insure that the ultimate argument of a derivation is valid whenever all gaps close. For, if all rules are sound, we can be sure that any interpretation that divides a gap and all its ancestors will divide some child and all ancestors of this child (since these are just the parent and its ancestors). But any interpretation that divides the ultimate argument of a derivation also divides any ancestor (since it has none), so if all rules are sound, this interpretation will also divide some child and all its ancestors—and so on. That is, as with strictness, when all rules are sound, an interpretation that divides the ultimate argument must divide some descendant at each stage; therefore, if all gaps close, there can be no interpretation dividing the ultimate argument. In short, if a sound rule ignores any gap-dividing interpretation, it is an interpretation that shows some risky conjecture does not follow from the initial premises, not one that shows that the initial conclusion was invalid.

Now, for a gap-closing rule to be sound, it is enough that there be no interpretation that makes the goal of the gap it closes false while making true all active resources of the gap *and all active resources of the gap's ancestors*. This means that it is enough for us to soundly close a gap that its goal be entailed by its active resources together the active resources of its ancestors. With the rules we have so far, all available resources are included if we take the active resources of a gap together with the active resources of its ancestors. So it is sound to close a gap when the goal is among the available resources, and our extended use of QED is sound.

But we can be even more generous since, by the law for lemmas, adding to a collection of resources something that is entailed by them will not change what they entail. In short, we can state rules for closing gaps and have them be sound if the conclusion of the gap is among its active resources, is among the active resources of its ancestors, *or is something entailed by these resources*.

The available resources of a gap always include its active resources and the active resources of its ancestors, but in 2.4.3 we will consider rules which add to the available resources certain conclusions entailed by these resources. And we have just seen that this sort of addition will not undermine the soundness of the extended use of QED.

Although we will sometimes need to distinguish soundness and safety (or even consider strictness) in later discussions, most often we will not. We will say that a system is *conservative* when its rules are all safe and sound (which, remember, comes to the same thing as being all safe and strict). So in a conservative system, gap-dividing interpretations are neither gained nor lost as we develop a derivation though they may be spread out among an increasing number of descendant gaps, something we will see illustrated in the next section's example.

2.3.5. Presenting counterexamples

A dead-end open gap is always divided by an interpretation, and any interpretation that divides it also divides the ultimate argument of the derivation. We will finish off derivations that uncover invalidity by displaying this division. We will do that by exhibiting an interpretation that divides a dead-end open gap and calculating the truth values of the original premises and conclusions on that interpretation. In the example discussed in 2.3.1 , this calculation is shown in the following table:

A	B	C	$(A \wedge \top) \wedge B \;/\; B \wedge (\top \wedge C)$
T	T	F	T T Ⓣ Ⓕ T F

Here the values of unanalyzed components have not been repeated on the right, but they are used to calculate the values of compounds containing them, with the order of calculation being guided by parentheses. In performing this calculation we are confirming that the interpretation dividing the gap really does constitute a counterexample to the ultimate argument; and we will say that, in constructing the table, we are *presenting a counterexample*. It will be our standard way of concluding the treatment of an argument whose derivation fails.

It is not always the case that all unanalyzed components of the ultimate argument all appear among the resources and goal of a dead-end gap. When unanalyzed components do not appear there, values must still be assigned to them in order for a truth value to be defined for each sentence in the ultimate argument; but it will not matter what value we assign to these further unanalyzed components. If an interpretation divides the gap, any way we choose to extend it to unanalyzed components not appearing in the gap's proximate argument will still divide that gap and therefore divide the ultimate argument.

The example below is designed to illustrate this. Of the three interpretations shown, the first divides only the first dead-end gap (since it assigns the value **T** to the goal of the second dead-end gap), and the last divides only the second open gap (for a similar reason); but the middle one divides both open gaps. With 4 unanalyzed components, there are $2{\times}2{\times}2{\times}2 = 2^4 = 16$ possible interpretations, so there are 13 interpretations that do not divide either gap. The soundness and safety of our rules insures that the 3 interpretations shown above constitute counterexamples to the ultimate argument and that the other 13 do not.

	$A \land B$	1
1 Ext	A	
1 Ext	B	(4)
	○	$A, B \nvDash C$
	C	2
	●	
4 QED	B	3
	○	$A, B \nvDash D$
	D	3
3 Cnj	$B \land D$	2
2 Cnj	$C \land (B \land D)$	

A	B	C	D	$A \land B$ /	$C \land (B \land D)$		
T	T	F	T	T	F	T	divides first dead-end gap
T	T	F	F	T	F	F	divides both dead-end gaps
T	T	T	F	T	F	F	divides second dead-end gap

While a dead-end gap is always divided by just one interpretation of the vocabulary appearing in its proximate argument, this interpretation may be provided by more than one interpretation of the derivation as a whole. That happens in both gaps here, and it also happens that a single interpretation of the whole derivation divides both of the gaps. That's why we end up with 3 interpretations all told.

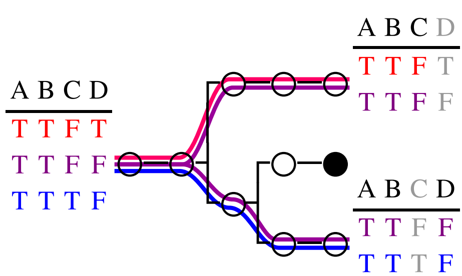

Fig. 2.3.5-1. The interpretations dividing the dead-end gaps of the example above.

Since each of these interpretations divides all ancestors of the dead-end gap or gaps that it divides, any one of the three is enough to provide a counterexample to the ultimate argument. Beginning with chapter 6, it will prove to be most convenient to assign **F** to an unanalyzed component whenever we have a choice, and here that would lead us to the middle interpretation in the case of both gaps. But, for now, when an unanalyzed component does not appear in the proximate argument of a dead-end gap, the choice of the value to assign to it is entirely arbitrary.

2.3.6. Reaching decisions

We know that if a system of derivations has individual rules that are both sound and safe and is, as a whole, sufficient, it will never give us an incorrect answer regarding the validity of an argument. But it is entirely possible that such a system will give us no answer at all. Of course, if we ever run out of rules to apply, we will have an answer. For then either all gaps will have closed or we will have an open gap that has reached a dead-end, and both results provide an answer. However, without some guarantee that we will eventually run out of rules, we have no guarantee that we will eventually have an answer. And such a guarantee is not trivial; in fact, once we get to the last two chapters, we will be working in a system some of whose derivations do go on forever.

We will say that a system is *decisive* when we always reach a point where either all gaps are closed or there is a dead-end open gap. It should be clear that our system so far is decisive. The rules Ext and Cnj replace conjunctions among the resources and goals of a gap by simpler sentences and must therefore eventually eliminate all conjunctions. And when the proximate argument of a gap contains no conjunctions, the only rules that might apply are QED, ENV, and EFQ. Each of these closes a gap and there will be only a limited number of gaps to close, so we must eventually run out of things to do.

But we will go on to consider further rules, and some of these will be sufficiently differently from those we have considered so far that, even when a system is decisive, it may not be as easy to see that it is. So let's look at some questions that arise in making this judgment. As we do this, it is worth remembering that, in assessing decisiveness, we are not really interested in whether a system reaches some valuable goal, only in whether we are bound to run out of things to do when we apply its rules.

One way to judge whether that is so is to provide some count of how much there is that might be done, and see whether each rule of the system reduces that count. However, it is not always easy to describe a single quantity that is always reduced, and the reason can be seen even with our current system. The rules QED, ENV, and EFQ reduce the number of open gaps, and that is certainly a relevant quantity. The rules Ext and Cnj, on the other hand, reduce the complexity of proximate arguments, something else that cannot go on for ever. While complexity may seem too abstract to be reduced to a single number, the simple expedient of counting the number of connectives in a proximate argument actually provides a useful quantity in the present setting. So far, so good, but the real problem arises in putting these two numbers together.

This problem is easiest to see by considering Cnj. While the proximate arguments of both its children are simpler than that of their parent, it adds to the total number of open gaps. It is tempting to say that this is acceptable because the increase in the number of open gaps is no greater than the decrease in the complexity, so the sum of the two is not increased. But this would be wrong on two counts. First, it is not enough that we avoid increasing the quantity we are watching: rules that merely kept it the same might go on for ever doing that. Second, our system would still be decisive if Cnj added 10, 100, or even a million new gaps when it eliminated a single connective. For, in the absence of a rule that added connectives, it would eventually run out of connectives to eliminate, and we would be forced to use other rules which did reduce one quantity without increasing the other.

This is not to say that there is no way of putting the number of open gaps and the complexity of proximate arguments together to produce a useful quantity, but any way of doing that must recognize their asymmetry: we can add gaps as we reduce the number of connectives but only provided we add no new connectives when we close gaps. However, we will not look at ways of actually combining these quantities. We will simply employ the abstract idea of a rule moving things along. We will call a rule that does this *progressive*, understanding that whether a rule is progressive depends not only on what quantities it might reduce but also on what other rules are present. The common idea associated with our various uses of this term progressive will be that, if all our rules are progressive, each moves us far enough along that we can never apply them more than a limited (though perhaps very large) number of times before we run out of things to do.

So a system all of whose rules are progressive will be decisive; that is, we will always reach a point at which no more rules can be applied. At that point, any gap that is left open will have reached a dead end, and the derivation will have provided an answer about the validity of the original system. And we saw earlier that if a system is sufficient and conservative, the existence or non-existence of an open gap when no more rules apply provides a correct answer regarding validity of the ultimate argument. A system that always eventually provides an answer and a correct one, can be said to provide a *decision procedure* for validity.

2.3.7. Soundness and completeness

Our current system is sufficient, conservative, and decisive, and it therefore provides a decision procedure. But we can cut up its properties in another way. Because it is decisive as well as accurate in its answers, we can say both of the following about any derivation:

(1) The ultimate argument of a derivation is valid if and only if at some stage all gaps have closed.
(2) The ultimate argument of a derivation is invalid if and only if eventually we reach a dead-end open gap.

The if parts of these together say that the system is accurate, and we have seen that they follow from its conservativeness (along with sufficiency in the case of the second statement). The only if parts follow from the if parts given decisiveness. (For example, if the ultimate argument is valid, it must be the case that all gaps close because otherwise, given decisiveness, we would reach a dead-end gap and the ultimate argument would not be valid.) Moreover, the only if parts of the two claims above together imply decisiveness because an argument will always be either valid or invalid, so they tell us that eventually either all gaps close or we reach a dead-end gap.

But these two claims, like the properties of soundness and safety, are not of equal importance. The first is closely tied to the use of derivations to establish validity while the second is similarly related to their use to find counterexamples and establish invalidity. The first is of special interest also because it can be established in some cases where decisiveness fails, and we will take it as the key property of our system of derivations in chapters 7 and 8 when we must abandon decisiveness.

It is standard to give different names to the two parts of the first statement:

(1a) The ultimate argument of a derivation is valid *if* at some stage all gaps have closed
(1b) The ultimate argument of a derivation is valid *only if* at some stage all gaps have closed

When we can be sure that (1a) is true, we say that the system is *sound*. We have seen that a system will be sound in this sense if all its rules are sound. When we can be sure that (1b) is true, we say the system is *complete* because such a system provides a proof for each valid argument.

We can show that a system is complete if we know (i) that its rules are safe and the system as whole is sufficient and we know also that (ii) any derivation whose ultimate argument is valid eventually reaches an end. Property (ii) is not full decisiveness since it applies only to derivations whose ultimate argument is valid. This sort of partial decisiveness is something we will be able to establish for the systems of chapters 7 and 8, for which full decisiveness does not hold. And, because this partial decisiveness is enough to provide completeness, all systems that we will study in the course are both sound and complete.

2.3.8. Formal validity

As was noted earlier, the use of the term valid in connection with derivations requires some qualification. In the context of derivations, as in the context of analyses, Roman capital letters are used to stand for particular sentences that are not analyzed further, and such sentences need not be logically independent. That means that a given extensional interpretation of unanalyzed sentences need not be realized in any possible world. So, in the example of 2.3.1, even though the appearance of a dead-end gap leads us to write "B, A, $\top \nvDash$ C", it might be that the particular sentences A and B do together entail the particular sentence C, and it could even be that C is tautology or that A and B are mutually exclusive. In short, knowing that there is an extensional interpretation of analyzed sentences that assigns them certain truth values does not show that it is logically possible for the sentences to have those truth values.

On the other hand, our interest in derivations is as a way of applying general principles of formal logic. And, even though these principles are applied to particular sentences, their application depends only on the features of these sentences that are displayed in symbolic analyses. In particular, the use of derivation rules does not depend on the specific identity of unanalyzed components. This means that when the gaps of a derivation do all close we know not only that its premises entail its conclusion but also that the same is true for any argument having the same form. One way of putting this is to say that we know the argument to be *formally valid* or, more precisely, to be valid in virtue of the form exhibited in the particular analysis we have used. Since formal validity is a stronger property than simple validity, knowing that an argument is formally valid is enough to tell us it is valid; and we will usually drop the qualification formal for this reason. But it is important to remember that when an argument is labeled "invalid" on the basis of a derivation, this judgment is relative to a particular analysis of it. Indeed, if this were not so, we could stop after studying conjunction: the point of considering further logical forms is to recognize the validity of arguments that count as formally invalid when considered solely in terms of conjunction.

The idea of validity in virtue of form can itself be spelled out by saying that an argument is formally valid with respect to a given analysis when any way of associating sentences with its unanalyzed components produces a valid argument. So when the derivation of 2.2.4 showed us that $(A \land B) \land C, D \vDash C \land (A \land D)$, this told us something not only about the specific sentences $(A \land B) \land C$, D, and $C \land (A \land D)$ but about any sentences that are related in the way indicated by these analyses—that is, about the sentences could be formed in these ways from any choice of sentences, A, B, C, and D. Such choice of actual sentences, one for each of a group of unanalyzed components, is an intensional interpretation in the sense discussed in 2.1.8, so we can say that an analyzed argument is formally valid when every intensional interpretation of it is valid.

When a derivation leads to a dead-end gap, what we know, speaking most strictly, is that its ultimate argument is not formally valid. That is because one test of formal validity is whether there is an extensional interpretation of the argument that divides its premises from it conclusion. And we will look more closely at why that is so.

First, if there is an extensional intepretation that divides an argument, we can construct an intensional interpretation by assigning to each component an actual sentence with the truth assigned by the extensional interpretation, and this interpretation will yield an actual argument having the same form as the original one but with actually true premises and an actually true conclusion. In example from 2.3.1, the counterexample given by the dead-end gap assigns **T** to A and B and **F** to C. So we might associate English sentences with these unanalyzed components as follows:

A: Atlanta is in Georgia
B: Boston is in Massachusetts
C: Chicago is in Massachusetts

If so, the proximate argument of the dead-end gap will be

Boston is in Massachusetts
Atlanta is in Georgia
$\top$

Chicago is in Massachusetts

and the ultimate argument of the derivation will be

Atlanta is in Georgia and $\top$; moreover, Boston is in Massachusetts

---

Boston and Chicago are both in Massachusetts

To get something completely in English, we can replace $\top$ by any tautology. If we use Atlanta is Atlanta, we get

Atlanta is in Georgia and is Atlanta; moreover, Boston is in Massachusetts

---

Boston and Chicago are both in Massachusetts

Each of these particular arguments has a false conclusion along with true premises not merely in *some* possible world but in the actual world, so they are certainly invalid. Because the latter two have the same form as the ultimate argument of the derivation, that ultimate argument is not valid *with respect to the form displayed in its analysis*. If in that argument, the unanalyzed A, B, and C happen to be sentences such that A, B $\vDash$ C, the argument will in fact be valid. For example, it might be

All humans are mortal and are human;
moreover, Socrates is human

Socrates is both human and mortal

But it will remain true that it is not valid with respect to the form displayed in the symbolic analysis, and we have shown it is not by giving another interpretation of this form that is not valid.

We have seen that an argument divided by an extensional interpretation is not formally valid. The converse is also true. That is, if an argument is not formally valid, its premises are divided from its conclusion by some extensional interpretation. The claim that an argument is formally valid is a generalization about both intensional interpretations and possible worlds, and a counterexample to this generalization is provided an intensional interpretation and a possible world with the property that the actual argument that results from the intensional interpretation is divided by the possible world. But any intensional interpretation and possible world will determine an assignment of truth values to the unanalyzed components of the argument. In the example above the value **T** is assigned to the unanalyzed component A by associating the sentence Atlanta is in Georgia with A and considering the truth value of this sentence in the actual world. Since any intensional interpretation and possible world will determine an extensional interpretation in this way, any counterexample to the formal validity of a symbolic argument will provide an extensional interpretation that divides its premises from its conclusion.

This means that even if we do not define formally validity directly in terms of indivisibility by extensional interpretations but instead in terms of validity under any intensional interpretation, it will still be true that an argument is formally valid if and only if no extensional interpretation divides its premises from its conclusion.

2.3.s. Summary

1 When a derivation is constructed for an invalid argument, we eventually reach a point where an open gap has reached a dead end without closing. We mark such a gap with a empty circle ○ and write its active resources and goal with the sign $\nvDash$ between to indicate that they do not form a valid argument. And we will see that the invalidity of the proximate argument of a dead-end gap implies the invalidity of the ultimate argument for which the derivation is constructed.

2 We will often be concerned with formal validity, so we extend to assignments of truth values the ideas of dividing premises from a conclusion and of constituting a counterexample to an argument. And we speak of a gap being divided when its proximate argument is. The fact that any dead-end open is divided—that its proximate argument has a counterexample—indicates that our system is sufficient in the sense of having enough rules to close all dead-end gaps whose proximate arguments are valid.

3 When speaking of the tree structure of the gaps of a proof, it is convenient to use a genealogical metaphor and to speak of a gap at one stage as the parent of the gaps that derive from it at the next stage, gaps that are its children. Children of a gap's children, their children, and so on are descendants of the gap, and it is an ancestor of them. We can state a necessary and sufficient condition for the divisbility of a gap in terms of the existence of divisible descendants at later stages.

4 We can be sure that a counterexample to the proximate argument of a dead-end gap is a counterexample to the derivation's ultimate argument provided all our rules are safe in the sense of never introducing new ways of dividing gaps. When the converse is true, when we our rules never allow us to ignore ways that a gap might be divided, they are strict. Since our real interest is in the ultimate argument of a derivation, it is really enough to attend to dividing intepretations when they also divide all ancestors of a gap. Rules that insure that we do this are sound; when all rules are safe, sound rules are also strict. The idea of soundness enables us to justify the use of available but inactive resources (to, for example, close gaps) even when not all rules are safe. A system whose rules are all sound and also safe is conservative.

5 When a dead-end open gap is divided by an interpretation, this

interpretation is also a counterexample to the ultimate argument of the derivation, and we will present such a counterexample as a way of finishing off a derivation that fails.

6 A system will be decisive (in the sense that any derivation will always come to an end) provided its rules are all progressive (in the sense of always leading us closer to a point where no more can be done). Many rules are progressive because they either close a gap or replace a goal or active resource by one or more simpler sentences. A decisive system which is sufficient and conservative (and is therefore correct in the answers it gives) provides a decision procedure for formal validity.

7 Not all systems we consider will provide decision procedures but all will be sound in the sense of providing proofs only for valid arguments and complete in the sense of leading us to a proof whenever an argument is formally valid.

8 An argument that is valid may have a form that is invalid in the sense that some intensional interpretation of the unanalyzed components appearing in the form—i.e., some way of associating actual sentences with them—yields an invalid argument. Formal validity implies validity, so a derivation that succeeds shows both, but one that fails only shows formal invalidity.

Glen Helman 03 Aug 2010

2.3.x. Exercise questions

Use the basic system of derivations to check each of the claims below; if a derivation indicates that a claim fails, present a counterexample (that is, give an interpretation that divides an open gap and calculate truth values for the premises and conclusion from it—as is done in the example in 2.3.5, though you need only provide a single counterexample even when the derivation leads you to several):

1. $A \vDash A \wedge B$

2. $A \wedge B \vDash A \wedge (B \wedge A)$

3. $B \wedge E, C \wedge \top \vDash (A \wedge B) \wedge (C \wedge D)$

4. $A \wedge B, B \wedge C, C \wedge D \vDash A \wedge D$

5. $A, B \wedge A, D \vDash B \wedge ((C \wedge A) \wedge D)$

For more exercises, use the exercise machine.

Glen Helman 03 Aug 2010

2.3.xa. Exercise answers

1.

	A	(2)
	●	
2 QED	A	1
	○	$A \not\models B$
	B	1
1 Cnj	$A \wedge B$	

A	B	A	/	A ∧ B
T	F	T		F

2.

	$A \wedge B$	1
1 Ext	A	(4),(6)
1 Ext	B	(5)
	●	
4 QED	A	2
	●	
5 QED	B	3
	●	
6 QED	A	3
3 Cnj	$B \wedge A$	2
2 Cnj	$A \wedge (B \wedge A)$	

3.

	$B \wedge E$	1
	$C \wedge \top$	2
1 Ext	B	(5)
1 Ext	E	
2 Ext	C	(7)
2 Ext	$\top$	
	○	$B, C, E, \top \nvDash A$
	A	4
	●	
5 QED	B	4
4 Cnj	$A \wedge B$	3
	●	
7 QED	C	6
	○	$B, C, E, \top \nvDash D$
	D	6
6 Cnj	$C \wedge D$	3
3 Cnj	$(A \wedge B) \wedge (C \wedge D)$	

A	B	C	D	E	$B \wedge E$,	$C \wedge$	$\top$ /	$(A \wedge B)$	$\wedge$	$(C \wedge D)$
F	T	T	F	T	T	T	T	F	F	F

This derivation could have been ended after stage 4 when the first open gap has reached a dead end. Often answers will show a derivation continued further than necessary in order to show how the further steps would have worked out. The counterexample presented here divides both dead-end gaps; there are others that divide one of the two. Notice that ⊤ is not assigned a value at the left of the table. Since its value is fixed by the stipulation that it is a tautology, a value need not and cannot be assigned to it as part of an extensional interpretation.

4.

	$A \wedge B$	1
	$B \wedge C$	2
	$B \wedge D$	3
1 Ext	A	(5)
1 Ext	B	
2 Ext	B	
2 Ext	C	
3 Ext	B	
3 Ext	D	(6)
	●	
5 QED	A	4
	●	
6 QED	D	4
4 Cnj	$A \wedge D$	

Clearly, there is redundancy in the active resources of the gaps after stage 3. Since both gaps close, the exploitation of the second premise at stage 2 is not necessary (though it would be necessary before any gap could reach a dead end). It would be possible to state rules so that the resource B was not repeated at stages 2 and 3, but such repetition does not ordinarily enlarge derivations significantly and makes it easier to check whether rules have been applied fully and correctly.

5.

	A	(6)
	$B \wedge A$	1
	D	(7)
1 Ext	B	(5)
1 Ext	A	
	· ●	
5 QED	· B	2
	· · · ○	$A,B,D \nvDash C$
	· · · C	4
	· · · ●	
6 QED	· · · A	4
4 Cnj	· · $C \wedge A$	3
	· · ●	
7 QED	· · D	3
3 Cnj	· $(C \wedge A) \wedge D$	2
2 Cnj	$B \wedge ((C \wedge A) \wedge D)$	

A	B	C	D	A,	$B \wedge A$,	D /	$B \wedge ((C \wedge A) \wedge D)$		
T	T	F	T	Ⓣ	Ⓣ	Ⓣ	Ⓕ	F	F

Glen Helman 03 Aug 2010

2.4. Using lemmas

2.4.0. Overview

Although our system of derivations as it stands is both sound and complete, we will add rules that reflect the use of lemmas, both because of the importance of lemmas in ordinary explicit deductive reasoning and because the sorts of organization and simplification they provde in that context are of value for our work, too.

2.4.1. The dangers of lemmas

Although the use of lemmas is valuable in general, not all individual lemmas are valuable: the uncontrolled use of lemmas can lead us into blind alleys or delay the progress of a derivation.

2.4.2. Lemmas for *reductio* arguments

A lemma that is entailed by our goal is safe (though not necessarily progressive); this means that any lemma is safe when the goal is $\perp$.

2.4.3. Attachment rules

Lemmas are certainly safe when we know we can prove them. We will use such lemmas to add to the available resources. The sentences added in this way may be more complex than those already present, so this use of lemmas can interfere with decisiveness.

2.4.1. The dangers of lemmas

A fully general rule for introducing lemmas was cited in 2.3.4 as an example of an unsafe rule because the resources of a gap might not entail the lemma even when the proximate argument of the gap is valid. Such a rule would also prevent a system from being decisive because it would always be possible to develop a gap further by introducing a lemma. However, as was noted earlier, more limited rules for introducing lemmas can be safe, and we will see later that they can also be progressive. In this section we will look at the problems posed by lemmas more closely before going on, in 2.4.2 and 2.4.3, to consider a couple of special cases where these problems do not arise.

The law for lemmas of 1.4.6 can be stated as follows:

$$\Gamma \vDash \varphi \text{ if both } \Gamma \vDash \psi \text{ and } \Gamma, \psi \vDash \varphi$$

Let us recall why this is true: any possible world that is a counterexample to the first entailment will be a counterexample to one of the two on the right—the first of them if it makes ψ false and the second if it makes it true. So if both entailments on the right hold (that is, neither has a counterexample), then the one on the left will hold, too.

But this principle is stated only as an if claim, and the corresponding only if statement is not always true. When $\Gamma \vDash \varphi$, we know that $\Gamma, \psi \vDash \varphi$ by the monotonicity of $\vDash$; but, since φ and ψ need have no connection with one another, knowing that $\Gamma \vDash \varphi$ would by itself give us no reason to suppose that $\Gamma \vDash \psi$. Of course, in a case where we know that $\varphi \vDash \psi$, we would know $\Gamma \vDash \psi$ because of the chain law, and there are other cases where would know $\Gamma \vDash \psi$ because of special connections between Γ and ψ. We will use lemmas in special cases like these; but, before turning to them, let's look at what a fully general rule for lemmas would be like.

If used in tree-form proofs a rule for lemmas would take the following form:

$$\text{Lem}\ \frac{\psi \qquad \begin{matrix}\cancel{\psi}\\[1.5em] \varphi\end{matrix}}{\varphi}$$

This rule divides the proof of φ into two branches, and represents a division of the task of proving φ into two components. The first is to prove the lemma ψ, and the second is to prove φ using the lemma ψ in addition to the already available assumptions and conclusions derived from them. The occurrence of ψ above φ at the right is intended to indicate that ψ is an assumption from which conclusions can be drawn in this branch, and the slash through it indicates that is not an assumption of the proof after the use of Lem. Such an assumption is said to be *discharged*, and we will use the same term when an assumption is no longer available in a derivation.

The assumption ψ may occur at the tips of several branches on the right, and it is legitimate to discharge some or all of them. The effect of Lem is understood to be the same as replacing the discharged assumptions ψ in the right branch with conclusions proved in the way shown by the left branch, so the effect of the rule is to erase the discharged assumptions as assumptions.

The value of using Lem rather than proving ψ on the right to begin with lies in allowing us to consider the two components of the argument separately and to save some work in cases where the assumption ψ appears more than once in the tree on the right. Both of these functions lie behind the use of lemmas in mathematical proofs. In the next section, we will see a simple example of the second sort of use.

The appearance of a corresponding step in a sequent proof is shown below. Notice that the assumption ψ is dropped between the second of the premise sequents and the conclusion sequent.

$$\text{law for lemmas}\ \frac{\Gamma \vDash \psi \qquad \Gamma, \psi \vDash \varphi}{\Gamma \vDash \varphi}$$

Sequent proofs with such a rule would work in only one direction. The law for lemmas does justify the conclusion sequent if we are able to establish the premise sequents, but we cannot go in the other direction. That is, since the conclusion sequent may hold even if the first of the premise sequents does not hold, we cannot investigate the requirements for an entailment to hold by moving up from the root.

In the notation of derivations, we use scope lines to mark the scope of added assumptions, which are marked off from other resources along a scope line by the sort of horizontal line we use to indicate the premises of the ultimate argument. In the diagram below, notice that the proximate arguments of the gaps before and after the rule is applied follow the pattern shown in the sequent proof step above.

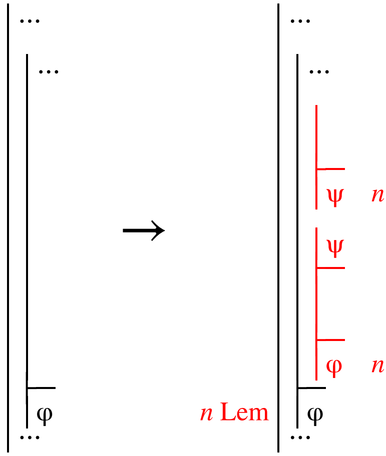

Fig. 2.4.1-1. Developing a derivation by introducing a lemma at stage *n* (a rule that will be part of our systems of derivations only in more restricted forms).

The assumption ψ is available only to the right of its scope line. After that scope line ends, it is said to have been discharged. The part of the proof in which this assumption is available corresponds to the right side of the tree in the tree-form and sequent proofs we have been looking at.

The effect of this rule on an argument tree is the following pattern of branching:

$$\frac{\Gamma \,/\, \psi \qquad \Gamma, \psi \,/\, \varphi}{\Gamma \,/\, \varphi}$$

Again notice that ψ is added to the premises in the right-hand child, and it is no longer among the assumptions when we move from this child back down the tree to its parent. (Nothing rules out the possibility that ψ already appears in Γ; this would be equivalent to having a tree-form proof in which the use of Lem discharges only some of the occurrences of the assumption ψ.)

The second of the two new gaps in a derivation developed using Lem should be, if anything, easier to close than the original gap because it has a further resource. This increased ease is the point of introducing a lemma. The price we pay for this is the need to close the first new gap also. If the lemma is properly chosen, that may also be easier than closing the original gap; but, because this rule is unsafe, we cannot be sure in general that the first new gap can be closed at all even if the original one could be closed eventually in other ways. Because of this, when lemmas are introduced in ordinary deductive reasoning we must be prepared to backtrack, to abandon the attempt to work by way of the lemma and look for another approach to the proof. The notation of derivations is not designed to incorporate backtracking, so we will use lemmas only in cases where we can be sure there will be no need to do that.

Indeed, we will not incorporate the rule Lem in the general form given here into our system of derivations. Instead, we will employ more specific rules that are based on the idea behind it. Even in cases where we can sure backtracking is not necessary, the introduction of lemmas could interfere with decisiveness if there were enough safe lemmas to keep introducing them forever. So our restrictions on the use of alternatives to Lem will be more severe than would be required merely to insure their safety.

2.4.2. Lemmas for *reductio* arguments

We have seen that one case where a lemma is bound to be safe is when it is entailed by the goal we seek. That is, we can state following principle:

If $\varphi \vDash \psi$, then $\Gamma \vDash \varphi$ if and only if both $\Gamma \vDash \psi$ and $\Gamma, \psi \vDash \varphi$

which tells us that, when $\varphi \vDash \psi$, it is not only sound but also safe to introduce a lemma ψ in a derivation whose goal is φ. It is the `only if` part of this principle—and, more specifically, the part that says $\Gamma \vDash \varphi$ only if $\Gamma \vDash \psi$—that requires the assumption that $\varphi \vDash \psi$.

In order to apply the idea of this principle in derivations, we can look for convenient ways of insuring that $\varphi \vDash \psi$. The obvious valid arguments of this form among those we have identified so far are EFQ and the two forms of Ext. Although EFQ will prove to be the more important, Ext is a better source of examples at the moment and we will consider it first. Here is a derivation which uses the rule Lem to introduce a lemma that is the result of applying left extraction to the final goal.

	$A \wedge B$	1
1 Ext	A	(5),(9)
1 Ext	B	(4)
	●	
4 QED	B	3
	●	
5 QED	A	3
3 Cnj	$B \wedge A$	2
	$B \wedge A$	(7),(10)
	●	
7 QED	$B \wedge A$	6
	●	
9 QED	A	8
	●	
10 QED	$B \wedge A$	8
8 Cnj	$A \wedge (B \wedge A)$	6
6 Cnj	$(B \wedge A) \wedge (A \wedge (B \wedge A))$	2
2 Lem	$(B \wedge A) \wedge (A \wedge (B \wedge A))$	

Here the rule Lem is applied at stage 2 with the left component of the goal as the lemma. This yields a slight shortening of the derivation since we are able to use the lemma to conclude B ∧ A by QED at stages 7 and 9 rather than repeating the proof used at stages 3-5 twice.

The basic idea here—isolating a component of an argument to avoid repeating it—is an important one. However, the actual simplification in this case is limited, and we would have few opportunities to use lemmas whose safety was assured by Ext. So we will not build this use of lemmas into our system of derivations, it will serve us only as an initial example.

The pattern *Ex Falso Quodlibet* provides the basis for a much more imporant use of lemmas. An argument whose conclusion is ⊥ is often called a *reductio argument*; reductio here is short for the Latin phrase *reductio ad absurdum* ('reduction to absurdity'). We will often need to use a lemma to complete such an argument and, since EFQ tells us that ⊥ entails any sentence, we know that any lemma we choose is safe. We will call the rule implementing this idea *Lemma for Reductio* or LFR:

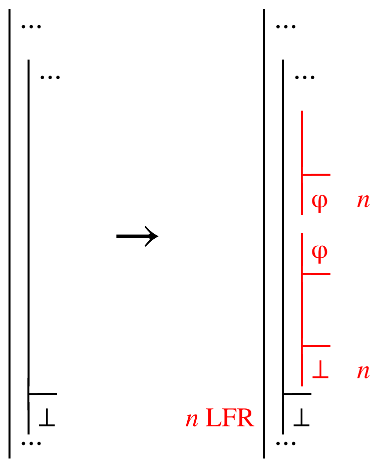

Fig. 2.4.2-1. Developing a derivation by introducing a lemma for a *reductio* at stage *n*.

The principle associated with this rule is the following:

Γ ⊨ ⊥ if and only if both Γ ⊨ φ and Γ, φ ⊨ ⊥

Although this follows from the principle stated at the beginning of the section and the validity of EFQ, it is instructive to look at its justification more directly. The if part again is just an instance of the law for lemmas. The only if part tells us that any interpretation dividing one of the child gaps will also divide the parent. But this must be so because any interpretation dividing a child will make the active resources of the parent true (since they all remain active in each of the children) and every interpretation makes ⊥ false.

Unfortunately, we are not yet in a position to illustrate this rule because we have no non-trivial examples of formally valid *reductio* arguments. A *reductio* is formally valid only if its premises constitute a formally inconsistent set (that is, one whose members cannot be all true on any extensional interpretation) and the only formally inconsistent sets available with our current analyses of sentences contain $\bot$ either as a member or as a component of one. Such a set a can be shown to entail $\bot$ with use of nothing but Ext and QED, so introducing LFR would merely complicate that argument.

In the next chapter, the rule LFR will serve us as a temporary expedient, but we will eventually introduce other special rules that are designed to cover the case where LFR would be most useful, and LFR itself will be ignored. One reason is that the free use of LFR would undermine decisiveness since the form of the rule places no constraints on the number of different lemmas that might be introduced. Something like a limitation to sentences that already appear as components of active resources and goals would be sufficient to insure decisiveness and would still permit the more important uses of the rule, but we will not attempt to formulate the sort of restriction that would enable us to prove decisiveness for a system with LFR. It is simply not important enough to bother. Apart from its role as a temporary expedient, it will serve us mainly as a way of displaying the connection between the special rules to be introduced later and the idea of a lemma.

2.4.3. Attachment rules

The second sort of case in which use of a lemma is bound to safe is one in which it is clearly entailed by the available resources. A principle justifying that use would take the following form:

If $\Gamma \vDash \psi$, then $\Gamma \vDash \varphi$ if and only if $\Gamma, \psi \vDash \varphi$

If we were to drop the only if, this would be just a another way of stating the law for lemmas, and we know that $\Gamma \vDash \varphi$ only if $\Gamma, \psi \vDash \varphi$ by monotonicity.

There might seem to be little point in regarding this as a case of a lemma at all. If a statement is entailed by the resources, why not just have a rule that allows us to add it to our active resources? Indeed, as far as soundness and safety are concerned, this would come to the same thing. But this second use of lemmas is motivated (as well as constrained) by considerations of decisiveness. That is, our system is decisive in part because active resources are added only to reduce the complexity of proximate arguments, and it may be useful to reach our goal by way of a lemma that is entailed by our resources but is more complex than they are. Adding such a sentence as an active resource would open the door to going around in circles in which we add complexity only to simplify and then add complexity again, and so on.

When discussing the soundness of QED in 2.3.4, we saw that it would be legitimate for a rule to close a gap when its goal is not among its active resources—or even among the active resources of its ancestor gaps—if the goal was entailed by available resources. We will not use such a sweeping rule but we will introduce a few rules in special cases that add to the available resources of a gap without changing either its active resources or its goal.

An example is the following way of developing a gap, which we will call *Adjunction*:

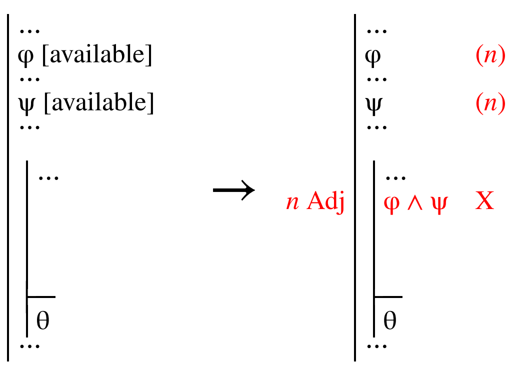

Fig. 2.4.3-1. Developing a derivation by applying Adj at stage *n*.

The added conjunction functions as a lemma, so this rule represents a way of using lemmas. However, it has a number of special features, both by comparison with a rule like LFR and by comparison with other rules we have seen.

The lemma φ ∧ ψ does not lie to the right of a new scope line, as it does in the second gap introduced by LFR, for two reasons. First, we have not branched the gap so the added resource is available throughout the gap. And, second, we do not need to mark this new resource off as an added assumption because it is entailed by those already present.

Notice also that we treat this rule not as a way to plan for our goal but simply as a way to add resources. However, it does not exploit resources in order to add others and the X to the right of φ ∧ ψ is intended to indicate that this resource need not be exploited further. One way to think about this is to suppose that φ ∧ ψ has been introduced as something already exploited. That is, although it need not have once been an active resource that has since been exploited (and it would already be part of the available resources if it had been), it has a status similar to such resources.

One example of the use of Adj is provided by the example in 2.4.2 (though we are thinking of the lemma differently now: there we thought of it as something entailed by the goal while here we think of it as something entailed by the resources).

	A ∧ B	1
1 Ext	A	(2), (6)
1 Ext	B	(2)
2 Adj	B ∧ A	X, (4), (7)
	●	
4 QED	B ∧ A	3
	●	
6 QED	A	5
	●	
7 QED	B ∧ A	5
5 Cnj	A ∧ (B ∧ A)	3
3 Cnj	(B ∧ A) ∧ (A ∧ (B ∧ A))	

With two more uses of Cnj, we would not have needed Adj; and, with two more uses of Adj, we would not have needed Cnj. Still, it is this sort of mixed use of the two rules that brings us closest to typical patterns of explicit deductive argument.

This example also exhibits the sort of foresight or insight that is required to Adj and similar rules. At stage 2, after simplifying our resources as far as possible, we look ahead to the analysis of the goal (before we have built that analysis into our derivation at stages 3 and 5), and we see that we will need to establish B ∧ A twice to reach it. Noticing that we already have the makings of this sentence among our resources, we then assemble it using Adj to have it available for later use.

Adjunction is one example of a group of rules we will refer to as *attachment rules*. Any such rule *R* will exhibit the following general pattern.

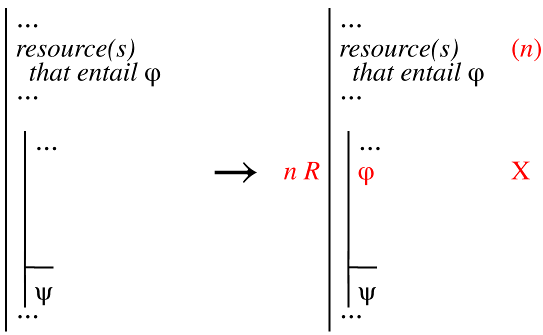

Fig. 2.4.3-2. Developing a derivation by applying an attachment rule *R* at stage *n*.

Since the lemma φ is not an active resource, the proximate argument of child gap is the same as the parent's proximate argument. This means that safety and soundness (even strictness) hold as they would for a gap that is completely unchanged. A rule like this must be considered when arguing for the soundness of rules like QED that use merely available resources, but the required argument was already considered in 2.3.4: an interpretation that divides a gap and all its ancestors will already make true not only all the available resources but also any sentence φ that is entailed by them, so we do nothing to change the situation by adding such a sentence φ to the available resources.

Of course, a rule like Adj does raise questions about decisiveness since the lemma it introduces is more complex than the premises it is based on. This increased complexity will be typical of attachment rules and is the reason for their name. We will not state the sort restriction on the use of attachment rules that would enable us to prove decisiveness; and, for practical purposes, the most valuable constraint on their use is simple good sense. But, as a rule of thumb, it is natural to limit the use of such rules to cases where the lemma is a component of a goal or active resource since such cases will represent the principal grounds for using these rules in any case. But it is important to remember that a sentence is a component of itself, and one common use of these rules will be to introduce the goal itself as an available resource in order to apply QED.

2.4.s. Summary

1 Using a lemma is one way of dividing up the work of a proof. We might use lemmas in derivations by dividing a gap into two gaps, one with the lemma as a goal and the other with it as a further assumption to use in reaching the original goal. A rule Lemma (Lem) that does this is not safe in general, and we will use only special instances of it.

2 A lemma is always safe when it is entailed by the current goal. We can use this idea in *reductio* arguments, arguments whose goal is ⊥. Since ⊥ entails any sentence, the rule that introduces lemmas in such circumstances, Lemma for *Reductio* (LFR), will be safe (though some restriction on its use is needed to insure it is progressive).

3 A lemma is also safe if we know we can establish it. We will use this sort of lemma only in attachment rules, rules that add the lemma as an available but inactive resource. The first example of this sort of rule is Adjunction (Adj) which adds a conjunction when both conjuncts are already available. Although attachment rules can help us to close gaps sooner, care is needed in their use if they are to be progressive.

The derivation rules we have so far are summarized in the table below. The names of the rules are links to the point in the text where they were initially described; look there to see the actual form taken by the rule.

Rules for developing gaps			*Rules for closing gaps*		
	for resources	*for goals*	*when to close*	*rule*	
conjunction $\varphi \wedge \psi$	Ext	Cnj	the goal is also a resource	QED	
			⊤ is the goal	ENV	
			⊥ is a resource	EFQ	*Basic system*
			Attachment rule		*Added rules (optional)*
			added resource	*rule*	
			$\varphi \wedge \psi$	Adj	
			Rule for lemmas		
			prerequisite	*rule*	
			the goal is ⊥	LFR	

2.4.x. Exercise questions

Use the basic system of derivations along with the attachment rule Adj to establish the following. These repeat entailments from earlier exercises and examples (specifically, **b** and **d** of exercise 2.2.x.2 and exercises **2** and **4** of 2.3.x). They will work best as exercises in the use of Adj if you avoid using Cnj.

1. $A \vDash A \land A$

2. $A \land B, B \land C, C \land D \vDash A \land D$

3. $A \land B \vDash A \land (B \land A)$

4. $A, B \land C, D \vDash (C \land (B \land A)) \land B$

The exercise machine doesn't incorporate attachment rules, so, while it can generate exercises where Adj would be useful, that rule won't be used in any answers it produces.

Glen Helman 03 Aug 2010

2.4.xa. Exercise answers

The answers below avoid the use of Cnj in order to maximize the use of the rule Adj. In some cases, a mixed use of the two would have produced a more natural argument.

1.

	A	(1)
1 Adj	$A \land A$	X,(2)
	●	
2 QED	$A \land A$	

2.

	$A \land B$	1
	$B \land C$	2
	$B \land D$	3
1 Ext	A	(4)
1 Ext	B	
2 Ext	B	
2 Ext	C	
3 Ext	B	
3 Ext	D	(4)
4 Adj	$A \land D$	X,(5)
	●	
5 QED	$A \land D$	

3.

	$A \land B$	1
1 Ext	A	(2),(3)
1 Ext	B	(2)
2 Adj	$B \land A$	X,(3)
3 Adj	$A \land (B \land A)$	X,(4)
	●	
4 QED	$A \land (B \land A)$	

4.

	A	(2)
	$B \land C$	1
	D	
1 Ext	B	(2),(4)
1 Ext	C	(3)
2 Adj	$B \land A$	X,(3)
3 Adj	$C \land (B \land A)$	X,(4)
4 Adj	$(C \land (B \land A)) \land B$	X,(5)
	●	
5 QED	$(C \land (B \land A)) \land B$	