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, unless there is some problem with our system, there must be derivations whose gaps do not all close. If the gaps of a derivation will never 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 so. We answered both questions in a preliminary way in 2.2.1 when considering tree-form proofs. We return to them now in order to consider the case of derivations more explicitly and to establish a framework for asking the same questions in later chapters. One byproduct of this discussion will be some ways of thinking about rules that will be useful when we consider some optional extra rules for derivations in the next section.

The short answer to the first of the 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. 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.

Here's a simple example of a derivation whose only open gap has reached a dead end.

	(A ∧ ⊤) ∧ B	1
1 Ext	A ∧ ⊤	6
1 Ext	B	(4)
	| •	
4 QED	| B	2
	| | •	
5 ENV	| | ⊤	3
6 Ext	| | A	
6 Ext	| | ⊤	
	| | ○	B, A, ⊤ ⇏ C
	| | C	3
3 Cnj	| ⊤ ∧ C	2
2 Cnj	B ∧ (⊤ ∧ C)	

This gap (marked with a ***white 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among the available resources of the gap; but they were exploited in the course of developing this gap, so they are no longer active. Finally, since the only active resources of the gap are B, A, and ⊤ and its goal is not ⊤ nor among the resources, we have no rule for closing the gap. In short, no rule of any of the three sorts can be applied at this point. We will use the white circle shown here to mark open gaps that have reached a dead end. And, also as is done here, we will write the sign ⇏ (***rightwards double arrow with stroke***) between the active resources and the goal. This indicates, roughly, that the active resources do not entail the goal; but its precise significance is discussed more fully below.

Notice that, while A ∧ ⊤ might have been exploited at any point after stage 1 in the derivation above, its components were not needed to close the other gaps. As a result, its exploitation can be postponed until stage 6. However, even though its components do not enable us to close the open gap, it must be exploited before the gap has reached a dead end. It is only after it has been exploited that there is no rule for developing the gap further.

Now, let's look more closely at what we can say in general about dead-end open gaps. First of all, such a 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 ⊤ as a goal or ⊥ 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 ⊤ and their goals are limited to unanalyzed components and ⊥; and no dead-end gap can contain an unanalyzed component both as an active resource and as its goal. And this means that we can assign truth-values to the unanalyzed components appearing in this 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 ⊤ and ⊥, the first can appear only as a resource and the second only as the goal so they will not interfere with having true resources and a false goal.

Such an assignment of truth values is an extensional interpretation in the sense defined in 2.1.7. In the case of the derivation above, an interpretation making the active resource of the dead-end gap true and its goal false is displayed in the table below. 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 (which mainly just repeat the assignments).

A	B	C	B,	A,	⊤	/	C
T	T	F	Ⓣ	Ⓣ	Ⓣ		Ⓕ

We will extend the use of the term *divide* that was introduced in 1.4.1 to describe what this sort of interpretation does. We will say that an extensional interpretation like this ***divides*** the active resources of a gap from its goal; and, when it does this, we will say that it divides the gap.

This terminology was originally introduced for arguments; and, in applying it here, we are thinking of the resources and goal of a gap as forming an argument. However, this is not the argument for which the derivation was originally constructed. 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. The argument formed from the active resources and goals of a dead-end open gap is the end of the line in this process. We will call the argument for which the derivation was originally constructed the ***ultimate argument*** of the derivation. When working on a particular gap, we are most immediately trying to show that the active resources of the gap entail its goal, so we are trying to show that the argument with these resources as premises and the goal as its conclusion is a valid one. We will call this argument the ***proximate argument*** of the gap. The proximate argument of a gap is "nearby" in the sense of being our immediate concern while our final goal is to decide whether the ultimate argument is valid. Notice that the ultimate argument of a derivation is the proximate argument of its initial gap.

We will refer to the extensional interpretation which divides the gap as a ***counterexample*** to the proximate argument of the gap. And, in writing B, A, $\top \not\Rightarrow$ C to the right of the gap, we say that the proximate argument of the gap is not valid. However, these references to counterexamples and invalidity require 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, in principle, such sentences need not be logically independent. That means that a given extensional interpretation of such sentences need not be realized in any possible world. So in the example above, it might be that the sentences A and B do together entail C and it could even be that C is tautology or that A or B is absurd. In short, knowing that there is an extensional interpretation of analyzed sentences that makes certain ones of them true and others false does not show that it is logically possible for the sentences to have these truth values.

On the other hand, our interest in derivations and tree-form proofs is as a way of applying general principles of entailment. And, even though these principles are applied to particular sentences, their application depends only on the features of these sentences that are displayed in the analysis of them that is shown by the symbolic notation. In particular, the use of rules does not depend on the specific identity of unanalyzed components. This means that when the gaps of a derivation 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 the argument is ***formally valid*** or, more precisely, is valid in virtue of the form exhibited in its analysis. 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. This sort of association of sentences with unanalyzed components is an intensional interpretation as defined in 2.1.7, so we can say that an analyzed argument is formally valid when every intensional interpretation of it is valid. We usually will not know the identity of the unanalyzed components of a symbolic argument, so formal validity is all that we will be in a position to judge; and we will often drop the qualification *formal*.

When a derivation fails, what we know, speaking most strictly, is that it's 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. If there is such a dividing interpretation, we can construct an intensional interpretation by assigning to each component an actual sentence with the truth assigned by the extensional interpretation, and this 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 above, if we associate sentences with unanalyzed components as follows:

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

we will have the invalid argument:

Boston is in Massachusetts
Atlanta is in Georgia
⊤
―――――――――――――
Chicago is in Indiana

which has a false conclusion along with true premises not merely in *some* possible world but, indeed, in the actual world. And, because this argument is invalid and has the same form as the proximate argument of the gap, the latter argument is not valid with respect to the form displayed in its analysis. This sort of thing will work with any example, so we know that, if the premises of an argument are divided from its conclusion by an extensional interpretation, the argument is not formally valid.

It is also true that, if an argument is not formally valid, its premises are divided from its conclusion by an extensional interpretation. A 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 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. So any intensional interpretation and possible world will determine an extensional interpretation, and any counterexample to the formal validity of a symbolic argument will provide an extensional interpretation that divides its premises from its conclusion.

So we will have an extensional interpretation dividing the premises and conclusion of an argument if and only if we have a counterexample to its formal validity. That means we can take formal validity to be a generalization about extensional interpretations: an argument is formally valid if and only if its conclusion is true under every extensional interpretation that makes its premises all true. This means that an extensional interpretation that divides the premises of an argument from its conclusion amounts to a counterexample to formal validity.

We saw earlier that any dead-end open gap provides us with this sort of counterexample to formal validity. And that tells us that our system of derivations has enough rules, for it tells us that we are able to develop or close a gap whenever its proximate argument is valid. And, if the proximate argument is not valid, we would not expect to move further towards the completion of a proof. We will indicate this sort of completeness 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 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.

Sufficiency is important, but there are further properties we might expect to hold of a good system of derivations. For example, we know that the proximate argument of a dead-end open gap is not valid; but that does not by itself show that the ultimate argument of a derivation with a dead-end gap will always be invalid, and testing the validity of the ultimate argument is the reason we construct the derivation. Moreover, sufficiency does not imply that we will even eventually reach a point where either all gaps close or there is a dead-end open gap; that is, a sufficient system might lead us to derivations that develop forever. In the next two subsections, we will see that our current system of derivations is well-behaved in both these respects.