7.5.2. The role of generalizations in entailment

The special features of the laws we will state for the universals can be traced to two sources. One is the analogy with conjunction we have just explored. The other is a pair of differences between what we have said about universals and what we may say about ordinary conjunctions. The first difference lies in the fact that the principles of entailment for universals must hold for all structures, so they cannot depend on special assumptions about the range **R** of reference values. This means, in particular, that the set of "components" of a universal (i.e., its instances in an expansion by **R**) must be left indefinite while an ordinary conjunction has a definite and, indeed, finite set of components. This would make universals difficult to deal with were it not for their second difference from conjunctions. The components of an ordinary conjunction can be any pair of sentences so they need have nothing in common, and we must consider them individually; but the instances of a universal all follow the same pattern, differing only in occurrences of a single term, so we can speak of them all together by speaking of this pattern.

We will develop laws for universals by taking certain laws for conjunctions as our model and modifying them to take account of the differences just outlined. The laws for conjunction we will work from are the following:

$$\phi \wedge \psi \Rightarrow \phi$$
$$\phi \wedge \psi \Rightarrow \psi$$
$$\Gamma \Rightarrow \phi \wedge \psi \text{ if and only if } \Gamma \Rightarrow \phi \text{ and } \Gamma \Rightarrow \psi.$$

The last should be no surprise since it is our basic law for conjunction as a conclusion but, although the first pair of principles are clearly associated with the rule of Extraction, they are less far reaching than our law for conjunction as a premise. Our use of these less general principles is itself due to the differences between universals and conjunction: the law for conjunction as a premise says we can replace a conjunction by its components, but there is no hope of doing anything like this for a universal since it has no one definite set of "components."

What laws can we state for unrestricted universals that are analogous to these laws for conjunction? The first pair of laws for conjunction noted above together say that a conjunction implies each of its components. The analogous claim about an unrestricted universal is that it implies each of its instances. This is a principle known as ***universal***

instantiation:

$$\forall x\ \theta x \Rightarrow \theta\tau \text{ for each term } \tau. \text{ (That is, } \forall x \ldots x \ldots \Rightarrow \ldots \tau \ldots.\text{)}$$

For example, the sentence *Everything is fine and dandy* implies the claim *The number 2 is fine and dandy* as well as other sentences of the form τ *is fine and dandy*.

The principle of universal instantiation is not quite what we will take as our account of the unrestricted universal as a premise. Universal instantiation can be used along with the law of lemmas to develop a derivation by adding any instance of a universal premise as a further resource. And it is that use of universal instantiation that we call our ***law for the unrestricted universal as a premise***: for any term τ, we have

$$\Gamma, \forall x\ \theta x \Rightarrow \Sigma \text{ if and only if } \Gamma, \forall x\ \theta x, \theta\tau \Rightarrow \Sigma.$$

That is, given a universal premise, we may add any instance as a further premise. Note that the instance is added as a *further* premise. We cannot drop the universal because we cannot expect its content to be exhausted by a single instance; *Everything is fine and dandy*, for example, has implications for things other than the number 2. As you might expect, our inability to drop the universal from the premises will some cause complications when we try to implement this law in derivations.

Now let us look at the role of an unrestricted universal as a conclusion. Here we have the law for conjunction as a conclusion to use as a model. We have to expect changes, though, because that law gives separate consideration to each of the two components of the conjunction and we cannot expect to do this for the instances of a universal. Still, the law for conjunction points us in the right direction: we should look for some way of connecting the validity of a universal conclusion with the validity of arguments having its instances as conclusions. A connection like this is used in geometric proofs when we begin by saying, for example, "Let *ABC* be a triangle," and then go on to use our conclusions concerning *ABC* to justify general conclusions about all triangles. That is, we sometimes establish universal claims by ***generalizing*** from particular instances of them.

Clearly not every generalization from a particular instance will be legitimate. Certain premises may entail *The Empire State Building is tall* without entailing *Every building is tall*. In a geometric argument concerning a triangle *ABC*, we limit the information that we may use about the instance that we are considering to what we may establish concerning any triangle. For example, we ignore the fact we are using a diagram that shows *ABC* as acute or obtuse, and we probably avoid drawing it as a right triangle or an isosceles triangle to begin with. These restrictions are sometimes expressed by saying that we are arguing about an *arbitrary* or an *arbitrarily chosen* triangle. The idea is that what you say about the triangle *ABC* should hold for a triangle chosen at random or even one chosen by your worst enemy. Let us call an argument like this a ***general argument*** since it argues for an instance in a way that will hold generally for values in the domain of a universal. The law we are looking for should say that an unrestricted universal is a valid conclusion from given premises if we can establish an instance of it by a general argument. But we need to make this more precise.

In particular, we need to say how we can recognize a general argument just by looking at the logical forms of the sentences it involves. If we were to give instructions for making a general argument about a triangle *ABC*, one thing we might say is that we should not use any special assumptions about *ABC*. If we are going to generalize about triangles, we may assume that *ABC* is a triangle but we should not assume that it is acute or obtuse. This is just another way of saying that we should not use special information about this triangle, but it suggests an idea we can apply to arguments when we know only their logical forms.

Since we are considering arguments for unrestricted universals, we must be able to generalize not just about triangles, or some other limited class, but about everything; and that means we should use no assumptions at all about the term from which we wish to generalize. So we can say this: if we wish to generalize from an instance $\theta\tau$ to a universal $\forall x\ \theta x$, the term τ should not appear in our assumptions. You may have noticed a couple of jumps here. Saying we have an assumption containing τ is different from saying we have used that assumption, and saying that τ appears in an assumption is different from saying that the assumption provides special information about τ. For example, *The number 2 is fine and dandy and so is everything else* mentions the number 2 without constituting a special assumption about it. Still, the requirement that the term from which we generalize not appear in the assumptions is easy to check and using it will not limit the entailments we can establish, only the terms we can use to establish them.

This requirement is enough to rule out many unwarranted generalizations but it does not exclude them all. To see why, suppose we are arguing from the assumption *Everything is like itself*. One conclusion we can draw is *Wabash is like Wabash* and, in doing so, we have certainly used no special assumptions about Wabash. But this conclusion says that Wabash has the property of being like Wabash and that makes it an instance of the generalization *Everything is like Wabash*. But generalizing to that conclusion is surely unwarranted. The problem with this argument is that even though the term *Wabash* stands in no special relation to the assumptions, it does stand in a special relation to the universal conclusion *Everything is like Wabash*. In particular, it plays a special role in the predicate that the conclusion claims to be universal. These considerations suggest a second requirement for a general argument: if we wish to generalize from an instance $\theta\tau$ to a universal $\forall x\ \theta x$, the term τ should not appear in our conclusion; that is, it should not appear in θ.

There remains only one sort of problem to consider. Suppose our assumption is *Everything has its bad side*. We can conclude *Wabash has its bad side*. But we cannot go on to conclude *Wabash has everything*. Now the instance from which this conclusion would generalize is an instance for the term *Wabash's bad side* and this term does not appear in either the assumption or the conclusion, so it satisfies both of the requirements we have imposed so far. We could handle cases like this by requiring that terms on which we generalize share no vocabulary with either the assumptions or the conclusion. That would take care of this case (since *Wabash's bad side* shares vocabulary with both) and it would be more than enough to insure that an argument is general. Indeed, it would be enough to require, of a compound term, that its main functor not appear in the assumptions or conclusion (so, in the example above, the real problem is the appearance of the functor λx (x*'s bad side*) in the premise and not the appearance of the term *Wabash* in the conclusion). However, it is easier simply to prohibit generalization on compound terms. Unanalyzed terms that satisfy the first two requirements clearly share no vocabulary with the assumptions or conclusion so, for those terms, the first two requirements are enough.

We are now ready to state our ***law for the unrestricted universal as a conclusion***: for any unanalyzed term a appearing in neither Γ nor $\forall x\ \theta x$, we have

$$\Gamma \Rightarrow \forall x\ \theta x \text{ if and only if } \Gamma \Rightarrow \theta a.$$

Let us say that an unanalyzed term appearing in neither the premises or conclusion of an argument is *parametric*, or a ***parameter***, for that argument. In this vocabulary, the law says that an argument with an unrestricted universal conclusion is valid if and only if the premises entail an instance of the universal for a parameter. When arguments are stated in English, phrases like *let* a *be arbitrary* or *let us choose* a *arbitrarily* function as commitments to use the term a as a parameter.

Can be we sure that there will always be a parametric term on hand when we need one? Clearly, we would be stymied if our premises contained every term in the language. That is not a practical concern, but it shows that if there is always to be an appropriate term available for even any finite set of premises, our language must contain infinitely many parametric terms. Can we be sure that it does? If we are working with an idealized model of English, we can just stipulate that is does. Parametric terms are needed only for the inner workings of a derivation and need never appear in the initial premises and conclusion, so they do not have to be already available in the language. But even in real English there seems to be no shortage. The letter X is certainly overworked, but mathematicians always seem able to find one more symbol, no matter how many they are already juggling. If our imaginations falter, we can always resort to X′, X″, X‴, and so on—or else, X_1, X_2, ...

There is a more crucial question about this law: is it really true? We have been adding restrictions to insure that generalization is warranted. Can we be sure that we have enough? To see that we do, note first that the *only-if* part of the law is no problem. It says that a universal cannot be a valid conclusion unless any instance for a parameter is also valid, and this must be so because the universal implies all its instances.

So, let us consider the *if* part. To establish it, it is easiest to show that failure of the entailment Γ ⇒ ∀x θx implies failure of the entailment Γ ⇒ θa when a is parametric. Now, for the first entailment to fail, it must be possible to find a reference value that serves as a counterexample to the universal ∀x θx in some case where each member of Γ is true. Since the parameter a is foreign to both the premises and the universal conclusion, it can be made to refer to such a counterexample without interfering with the truth values of the premises. But this is to say we can make θa false while all members of Γ are true, and that means that the entailment Γ ⇒ θa fails. Turning this around, if Γ ⇒ θa holds when a is a parameter, then, when the premises Γ are all true, we know that there is no counterexample to the universal, which implies that the universal is true.

There is a more concrete way of showing that a parametric term is enough to insure that we have a general argument. Suppose we have established a conclusion θa for a parametric term a. Generalization to ∀x θx will be legitimate if we can argue from the same premises to each of its instances θτ. But we can find an argument for an instance θτ by following the pattern set by the argument for θa; we can simply replace the term a by τ everywhere in the argument for θa to get an argument for θτ. Because a is a parameter, it is not in the premises (and thus shares no vocabulary with them), so there is no connection between a and the premises that would make the argument go through for it but not for another term τ. And, because the parameter a is not in the conclusion, it is not in θ, so replacing it by τ everywhere will leave θ unchanged and change θa into θτ.

This argument recalls the comparison of the universal with conjunction. Since a conjunction can have any components, we must argue for each component individually and, since a conjunction has only two components, there is nothing to keep us from doing this. On the other hand, there would be no hope of providing a separate argument for each instance of a universal since, in general, there is no way of setting a limit on the number of instances it has. However, there is no need to consider each of these instances individually since they all have the same form, so an argument for one parametric instance can set the pattern for all of the rest.

Let us collect the two laws for the unrestricted universal before going on to see how to implement them in derivations.

> ***Law for the unrestricted universal as a premise***. For any term τ, we have:
>
> $$\Gamma,\ \forall x\ \theta x \Rightarrow \Sigma \text{ if and only if } \Gamma,\ \forall x\ \theta x,\ \theta\tau \Rightarrow \Sigma.$$
>
> ***Law for the unrestricted universal as a conclusion***. For any unanalyzed term a appearing in neither Γ nor ∀x θx, we have:
>
> $$\Gamma \Rightarrow \forall x\ \theta x \text{ if and only if } \Gamma \Rightarrow \theta a$$

Again notice the difference between range of instances mentioned in the two laws. In the first, τ can be any term while, in the second, a must be a parameter relative to the premises and the universal.