Invalidity in Predicate Logic

So far we've got a method for establishing that a predicate logic argument is *valid*: do a derivation. But we've got no method for establishing invalidity.

In propositional logic we had a method for establishing that an argument was invalid: that of truth tables. Recall what that method was exactly: we looked to see if there was a possible case (i.e. assignment of truth values to sentence letters) in which the premises were all true, but in which the conclusion was false. Notice that it just takes one case in which the premises are true but the conclusion false to make an argument invalid.

In predicate logic, too, this will be how we show that an argument is invalid: all we need to do is show that there's at least one case in which the premises are true but the conclusion false. For the whole idea of validity is that in a valid argument, it is *impossible* for the premises to be true while the conclusion is false. But we'll need to have a different notion of a "case". In propositional logic a case was simply an assignment of truth values to the sentence letters. But for a predicate logic formula like ∃xFx, we can't do this, because the formula "Fx" doesn't have a truth value on its own; rather, intuitively, it has different truth values depending on what x is. So the goal of this section is to come up with a new understanding of what a "case" is for predicate logic; the new kind of case we'll call a *model*.

Before we define what a model is, though, let's start with an example. The following argument is clearly invalid:

∃xFx
∃xGx

∃x(Fx&Gx)

The intuitive reason for this is as follows. The first premise says that something is an F. The second says that something is a G. But these "somethings" needn't be the same. Suppose that the class of Fs is completely disjoint from the class of Gs – no F is a G, and no G is an F. Then there will be some Fs, and some Gs, but nothing that is both.

Let's get specific. Imagine that there are only two people in the world; let's call the first person "0" and the second person "1". This we represent by the following:

U = {0,1}

"U" stands for the "universe" -- we let U be the set of all the people in the world that we're imagining.

Next, let's say something about what the predicates F and G apply to in our world. We do this by, in each case, writing down the set containing all the inhabitants of the world that the predicate applies to:

F: {0}
G: {1}

This is called specifying the *extension* of the predicates. By saying that the extension of the predicate F is {0}, we've said that the predicate F applies to only one thing in our imaginary world: 0. Since we didn't include 1 in the extension of F, this means that the predicate F *doesn't* apply to 1, in our imaginary world. Similarly, by letting the extension of G be {1}, we've said that the predicate G applies to 1, but not to 0. If F stood for "is female", and G stood for "is male", then our imaginary universe would consist of just two persons, 0, who is a female, and 1, who is a male.

This collection of i) a universe, and ii) a specification of the extensions of predicates, is called a *model*. Thus, our model can be displayed as follows:

U = {0,1}
F: {0}
G: {1}

It should be intuitively clear that the premises of our argument are "true in" this model, whereas the conclusion is false. The first premise says that there is at least one F, and there is, since 0 is in the extension of F. The second premise says that there is at least one G, and there is, since 1 is in the extension of G. The conclusion says that there is at least one thing that is *both* F and G, and this is false, because nothing is in *both* the extension of F, and also the extension of G. But we need more than this: we need a proof that the premises are true and the conclusion is false. This we do by the "method of expansions".

Consider the first premise, ∃xFx. This says that something is F. Now, since there are just two objects in the universe of our model, 0 and 1, this is equivalent to saying that *either* 0 is F, or 1 is F. Thus, in our model, ∃xFx is equivalent to:

F0 ∨ F1

Now we can compute the truth value of this formula by using the truth tables from before, as follows:

F0 ∨ F1

T   F

   T

We know F0 is true because 0 is in the extension of F. We know that F1 is false because 1 is *not* in the extension of F. And finally, we know that the truth value of the whole is T because T∨F yields T, from the truth table for ∨.

Similarly, we can calculate the truth values of the expansion of the second premise:

$$G0 \vee G1$$

F T

T

And finally we must calculate the truth value of the expansion of the conclusion. This is a bit trickier, because the conclusion is a more complex formula than the premises, but the principle is the same: $\exists x(Fx \& Gx)$ in our model means the same thing as: *either* 0 is both F and G, or 1 is both F and G. Thus, the expansion for $\exists x(Fx \& Gx)$ in our model is: $(F0 \& F1) \vee (G0 \& G1)$. In general, the rule for expanding existentials is:

> The expansion of a formula, $\exists x\phi[x]$ is "$\phi[0] \vee \phi[1] \vee \ldots$", where we have one disjunct for each member of the universe

OK, given this we can calculate the truth value of the expansion of $\exists x(Fx \& Gx)$ in our model:

$$(F0 \& F1) \vee (G0 \& G1)$$

T F F T

F F

F

This is the desired result: the conclusion is false in this model. So we've come up with at least one case -- i.e., at least one model -- in which the premises of the argument are true, but the conclusion is false. So the argument is invalid.

The steps, then, in showing an argument invalid are as follows:

1. Produce a model. This consists of specifying:
   a) a universe, and
   b) extensions

2. Produce the expansions of the premises and conclusion

3. Show via calculation that the expansions of the premises are true, whereas the expansion of the conclusion is false

Let's do another example, and this time I'll show how to do the expansion for a universal quantifier:

$\forall$x(Fx$\rightarrow$Gx)
$\exists$xGx & $\exists$x~Gx

$\exists$xFx

OK, how should we go about coming up with the model? The best rule of thumb is to *do forced steps first*. For example, the second premise forces us to do two things: add two objects to the universe, and put only one of them in the extension of "G". So we may begin with:

U = {0,1
G: {0
F: {

In general, if we have an existential to make true, it's often good to do this first, because it forces us to add something to the universe. Now, what other steps are we forced to do? Well, we want the conclusion to be false, and thus we can't have any object in the extension of "F". Therefore, we've got the following:

U = {0,1
G: {0
F: {}

(Notice that the extension of "F" is now "the empty set" -- this means that "F" doesn't apply to anything in our model.) OK; let's now ask whether we need to do anything else to get the first premise true. The answer is no. For the first premise says that "Fx$\rightarrow$Gx" is true no matter what x is. But since the predicate "F" applies to nothing in our model, "Fx$\rightarrow$Gx" will turn out vaccuously true no matter what x is (since its antecedant will always be false). Thus, our final model is:

U = {0,1}
G: {0}
F: {}

Now for the expansions. The first premise is a universal, and the rule for expanding for a universal is this:

> The expansion of a formula, $\forall x\phi[x]$ is "$\phi[0]$ & $\phi[1]$ & ...", where we have one conjunct for each member of the universe

Whereas existentials expand to disjunctions, universals expand to conjunctions. And this is the intuitively correct rule; $\forall$xFx says, for example, that *everything* is F, so if the universe consisted of just {0,1}, then $\forall$xFx would say that 0 is F *and* 1 is F -- that is, "F0 & F1". Applied to the first premise of the current argument, the expansion is:

```
(F0→G0) & (F1→G1)
 F   T      F   F
```

T T
T

Now, the expansions for the other premise and the conclusion:

$(G0 \vee G1) \& (\sim G0 \vee \sim G1)$
T F T F
F T
T T
T

$F0 \vee F1$
F F
F

Thus, our model does indeed establish that the argument is invalid, because the premises came out true whereas the conclusion came out false.

Our next example involves a name. To give a model, we'll need to know what the extension of a name is: the extension of a name is a *particular* member of the universe, *not* a set of things. That's because a name *refers to* a particular thing, unlike a predicate, which may *apply to* many things. Here's the example:

$\forall x(Fx \rightarrow Gx)$
Ga

Fa

Pretty clearly, the argument is invalid: the first premise says that all Fs are Gs, but this doesn't imply that all Gs are Fs. So a could be a G, but not an F. Here's a model based on this intuition:

U = {0}
F: {}
G: {0}
a: 0

Notice that I gave "a" the extension 0 -- in other words, "a" *refers to* 0. Now, for the expansions. The first premise is a universal, so the expansion should be a conjunction. But since there's only one member of the universe, the expansion will have only one "conjunct", and so there won't be any & (and so it won't really be properly called a conjunction). Thus, the expansion of premise 1 is:

$F0 \rightarrow G0$
F T
T

As for the second premise, to expand a formula containing a name, we simply replace that name with its referent:

G0
T

And finally, the conclusion:

F0
F

Good -- the premises came out true but the conclusion came out false.

Try $\forall xFx \rightarrow \forall xGx$ / $\forall x(Fx \rightarrow Gx)$ on your own.

The next example will illustrate how we give the extensions for *two* place predicates:

Lab

$\forall x Lax$

This is clearly an invalid argument: from the fact that a loves b, it doesn't follow that a loves *everyone*. a might love b, but not love him/herself, for example; if so, then since a doesn't love him/herself, a doesn't love everyone. So we can begin to construct the model as follows:

U = {0,1}
a: 0
b: 1

But now we want to say that 0 loves 1, but 0 does not love 0. We do this by specifying the following extension for "L":

L: {<0,1>}

The rule is this: *the extension for a two-place predicate is a set of <u>ordered pairs</u>*. We put in the extension of "L" the set of all and only ordered pairs, <i,j>, such that i L's j. Supposing "L" to stand for "loves", the extension of "L" would be the set of all ordered pairs where the first member loves the second.

We need to use pairs because a person doesn't just plain love -- rather, one person loves another. We need to use *ordered* pairs because, for example, a person might love another person without getting loved back. So, if the first person is 0 and the second is 1, we would put <0,1> in the extension of "L", but we wouldn't put <1,0> in.

Thus, our model is:

U = {0,1}
a: 0
b: 1
L: {<0,1>}

OK, now let's do the expansions. The expansion of the premise is simply:

L01

We simply put in the referents of "a" and "b". And to tell whether this is true, we look to see whether <0,1> is in the extension of "L". Since it is, we have:

L01
T

Now for the conclusion:

L00 & L01

This is the expansion because i) we simply stuck in 0 for "a", since 0 is the referent of "a", and ii) because of the universal quantifier ∀x, for each member of the universe, we included a conjunct where "x" referred to that member of the universe. The truth values are:

```
L00 & L01
F      T
    F
```

(L00 is false because the ordered pair <0,0> is not in the extension of "L".)

Next example:

∀x∃yLxy

∃y∀xLxy

The premise could be interpreted as meaning: everyone loves someone (or other); the conclusion would then mean: there is someone that is loved by everyone. Clearly the first doesn't imply the second, because the first doesn't imply that the person getting loved is the same in each case. So if we had a model with two things, in which each loved the other, but in which neither loved him or herself, then this would do the trick:

U = {0,1}
L: {<0,1>,<1,0>}

Now for the expansions. These are trickier because we have two nested quantifiers. We do this in parts; let's do the premise first. Think of $\forall x\exists yLxy$ as having the form:

$\forall x\ (\exists yLxy)$

Since this is a universal, we need to expand it to a conjunction:

$\exists yL0y\ \&\ \exists yL1y$

What we did was this: we took the inside part of the premise, $\exists yLxy$, and for each member of the universe, we included a conjunct where we put that member of the universe in for "x". Now we need to expand these two conjuncts. Since each contains the quantifier $\exists y$, each will turn into a disjunction:

```
(L00 ∨ L01) & (L10 ∨ L11)
  F      T        T       F
      T                T
             T
```

As for the conclusion, we do it in steps as well:

$\forall xLx0 \vee \forall xLx1$

```
(L00 & L10) ∨ (L01 & L11)
  F     T        T       F
      F                F
              F
```

Finally, we need to learn how to construct models for arguments involving the identity predicate; for example:

$\exists xFx$

$\exists x\exists y(Fx\&Fy\&x\neq y)$

In English, this is saying "There is at least one F; therefore, there are at least two Fs". This is clearly an invalid argument -- what if there is only one F? Thus, we choose the following model:

U: {0}
F: {0}

We have only one thing in the universe, and that thing is F. We now expand the premise and conclusion, noting that even though we have existentials, since there is only one member of the universe, the "disjunctions" we get aren't really disjunctions. Here is the premise:

F0
T

And here is the conclusion, which we do by parts, since it contains two quantifiers:

$\exists x \exists y(Fx \,\&\, Fy \,\&\, x \neq y)$

becomes

$\exists y(F0 \,\&\, Fy \,\&\, 0 \neq y)$

(Notice again that there's only one "disjunct" since there's only one thing in the universe; if there were two then this would be a disjunction containing two disjuncts.) This in turn becomes:

$F0 \,\&\, F0 \,\&\, 0 \neq 0$

Now, the truth values for the first two conjuncts are:

$F0 \,\&\, F0 \,\&\, 0 \neq 0$
T T

since 0 is in the extension of "F". But the final conjunct is false, because 0 *is* identical to 0 (i.e., 0 is identical to itself). In general, the rule for computing truth values of identities is this: identities involving the same member of the domain twice, like 0=0, 1=1, etc., are true; all others, like 0=1, 0=2, 1=2, etc., are false. Thus, the conclusion is false, since it has a false conjunct:

$F0 \,\&\, F0 \,\&\, 0 \neq 0$
T T F

\ /

F

Infinite model:

$\forall x \exists y Rxy$
$\forall x \forall y \forall z[(Rxy \& Ryz) \rightarrow Rxz]$

$\exists x Rxx$