Chapter 12: Methods of Proof for Quantifiers

§ 12.1 Valid quantifier steps

The two simplest rules are the elimination rule for the universal quantifier and the introduction rule for the existential quantifier.

Universal elimination

This rule is sometimes called **universal instantiation**. Given a universal generalization (an ∀ sentence), the rule allows you to infer any instance of that generalization.

Example:

From *Everyone is mortal*, infer *Dick Cheney is mortal*.

The formal version of this rule (to be developed in Chapter 13) is called ∀ **Elim**. It will permit inferences like the following.

Example:

From ∀x Cube(x), infer Cube(b).

Notice that the term "elimination" is somewhat misleading here, for nothing is really being eliminated. But since the sentence *from which* the inference is drawn contains a universal quantifier that does not occur in the sentence which is inferred from it, one might well think of this maneuver as "eliminating" the universal quantifier.

Clearly the inferences above are valid. There is no way the *from* sentence can be true while the *to* sentence is false. (We are assuming, in both cases, that the names being used denote objects in the domain of discourse.) If Dick Cheney is not mortal, then it is not true that *everyone* is mortal. And if Cube(b) is false, then we have a counterexample to ∀x Cube(x).

Existential introduction

This rule, which permits you to introduce an existential quantifier, is sometimes called **existential generalization**. It allows you to infer an existential generalization (an ∃ sentence) from any instance of that generalization.

Example:

From *Dick Cheney is mortal* infer *Someone is mortal*.

The formal version of this rule is called ∃ **Intro**. It will permit inferences like the following.

Example:

From Cube(b), infer ∃x Cube(x).

Again, these are both valid inferences. There is no way the *from* sentence can be true while the *to* sentence is false. If it is true that Dick Cheney is mortal, then it is true that someone (Dick Cheney, at the very least) is mortal. And if ∃x Cube(x) is false, then there are no cubes, and so Cube(b) is false.

> We thus have our first two simple rules for quantifiers: we can infer **from** a universal generalization **to** any of its instances, and we can infer **to** an existential generalization **from** any of its instances.

Why the next two rules are more complicated

We now have an elimination rule for the universal quantifier, and an introduction rule for the existential quantifier. This means that we can draw inferences **from** universally quantified sentences (∀ sentences), and **to** existentially quantified sentences (∃ sentences). Further, these rules are very simple, as can be seen from the examples above, and can be very simply stated (as can their formal counterparts ∀ **Elim** and ∃ **Intro**).

We also need to draw inferences **from** existentially quantified sentences and **to** universally quantified sentences. The question is, how can we formulate inference rules that enable us to do this?

It is clear that if we model these new rules on our present ones, they will not work. Such overly simple "rules" would look like this:

"Existential Elimination"

> From an existential generalization, infer any of its instances.

"Universal Generalization"

> From any instance of a universal generalization, infer that generalization.

The trouble is, inferences made in accordance with these "rules" are not valid, for they would permit us to infer false sentences from true ones. Here are some examples of **invalid** arguments that these "rules" would permit:

> Someone is a liberal
> ___
> Dick Cheney is a liberal

> ∃x Cube(x)
> ___
> Cube(b)

> Bill Bradley is a liberal
> ___
> Everyone is a liberal

> Cube(c)
> ___
> ∀x Cube(x)

It's obvious that these arguments are invalid. Bill Bradley is a liberal, but not everyone is (e.g., Dick Cheney isn't a liberal). Someone is a liberal (e.g., Bill Bradley), but Dick Cheney isn't. And it's easy to construct a Tarski's World counterexample to the other two arguments: let *b* be a tetrahedron and *c* a cube.

So our simple versions of the new rules will not do. We must therefore come up with different rules for drawing inferences **from** existential generalizations and **to** universal generalizations. Instead of introducing those rules at this point, we will informally describe a method of drawing an inference **from** an existential generalization, and a method of inferring **to** a universal generalization. Having done that, we will be in a position to formulate sound rules of existential elimination and universal introduction.

         Revised 11/23/04

§ 12.2 The method of existential instantiation

The method

We give up the idea of trying to **infer** an instance of an existential generalization from the generalization. Instead, we temporarily introduce a **new** name into our proof and **assume** that it names an object (whatever it might be) that makes the existential generalization true. (We know that there must be such an object—we just don't know its name.) Then we try to prove something about the hypothetical individual. Finally, we derive some further sentence (typically, an existential generalization) that does not mention that individual by name.

Example

Argument

Some criminal stole the diamonds from the museum. Whoever stole the diamonds has an accomplice on the museum staff. Therefore, some criminal has an accomplice on the museum staff.

Proof

We know that some criminal stole the diamonds; let's call him (or her) *Ralph*. Since whoever stole the diamonds has an accomplice on the museum staff, it follows that Ralph has such an accomplice. But Ralph is a criminal, and Ralph has an accomplice on the staff. So, some criminal has an accomplice on the staff.

The rule

If we have followed this method successfully, we are in the following situation:

- We have an existential generalization as a line in our proof, say $\exists x\, P(x)$.

- We have assumed an instance of that generalization, say $P(c)$, as a temporary assumption.

- We have derived from that assumption some conclusion, say Q, in which c does not occur.

The rule then permits us to enter the conclusion Q that we just reached as a new line, but one which depends on the existential generalization $\exists x\, P(x)$, rather than on the instance $P(c)$ we temporarily assumed.

Our example followed this procedure: $P(x)$ was *x is a criminal and x stole the diamonds from the museum*, c was *Ralph*, and Q was *Some criminal has an accomplice on the staff*. Our assumption came at the point where we said *Let's call him Ralph*.

The example on p. 323 also makes clear how this works. When we get to Ch. 13, we'll see how the rule for system $\mathcal{F}$ formalizes this procedure.

§ 12.3 The method of general conditional proof

Once again, we give up the idea of trying to **infer** a universal generalization from just any instance of the generalization. Instead, we temporarily introduce a **new** name into our proof and **assume** that it names an object chosen at random. Then we prove something about the randomly chosen individual. Finally, we may then infer that what we have proven about this randomly chosen individual holds universally—i.e., we may infer a universal generalization.

The method

This is a method of proving **generalized conditional** sentences, that is, sentences of the form *All P's are Q's*. The technique is to pick some arbitrary instance of *P*, and then prove that it is also an instance of *Q*. Having shown that this arbitrary instance of *P* is also an instance of *Q*, we may infer that **every** instance of *P* is an instance of *Q*.

One might well wonder, how can we be certain that we have picked an instance of *P*? What happens if there are none? The answer is that we do not have to be certain of this. That there is an instance of *P* (chosen by us at random) is just an assumption we are making, an assumption that we will discharge. So, in the end, our proof will not depend on there actually being such an instance. Rather, what we show is that **if** there is such an instance, it will also be an instance of *Q*. The method looks a lot like the conditional proof method we used in propositional logic.

That is, to prove a statement of the form $\forall x\ (P(x) \rightarrow Q(x))$:

- Assume some instance of the wff $P(x)$, say $P(c)$, where c denotes any arbitrarily selected individual satisfying $P(x)$.

- Prove $Q(c)$

- Discharge the assumption and draw the conclusion $\forall x\ (P(x) \rightarrow Q(x))$.

There is another way to look at this kind of proof, one that usually goes by the name **universal generalization**. Here, one starts out with only the assumption that one has chosen some object at random (but no other assumption about it). One then proves something about this object. One then concludes that whatever one has proved about this arbitrarily chosen object holds of every object. That is:

- Choose a name, say c, and assume that it denotes some arbitrary individual.

- Prove some sentence containing the name c, say $S(c)$.

- Discharge the assumption and infer the universal generalization $\forall x\ S(x)$.

As *LPL* points out, these two approaches are in fact redundant—we can make do with only one of them. But the first is common in everyday reasoning, and the second is common in logic books, so we will build them both into system $\mathcal{F}$ and into Fitch.

Planning a strategy: informal proofs

Sketching out an informal proof is almost always a good thing to do before trying to construct a formal proof. So before moving on to the next chapter, let's try our hand at some informal proofs.

Example: Exercise 12.9

$\forall x\ [(Cube(x) \wedge Large(x)) \vee (Tet(x) \wedge Small(x))]$

$\forall x\ [Tet(x) \rightarrow BackOf(x, c)]$

$\forall x\ [Small(x) \rightarrow BackOf(x, c)]$

Proof: We will use the method of **general conditional proof**. Let *a* be an arbitrary small object—we will prove that *a* is back of *c*. Premise 1 tells us that every object is either a large cube or a small tetrahedron, so it follows (by $\forall$ **Elim**) that *a* is either a large cube or a small tetrahedron. This gives us two cases. But the first case is immediately contradictory, since *a* cannot be both small and large, so the second case must hold, and *a* must be a small

tetrahedron. But premise 2 tells us that every tetrahedron is back of *c*, so it follows that *a* is back of *c*, which is what we wanted to prove. Our assumption that *a* is small, then, has led to the conclusion that *a* is back of *c*. Hence, by general conditional proof, anything that is small is in back of *c*. QED.

Example: Exercise 12.20

> ∀x [Cube(x) → ∃y LeftOf(x, y)]
>
> ¬∃x ∃z [Cube(x) ∧ Cube(z) ∧ LeftOf(x, z)]
>
> ∃x ∃y [Cube(x) ∧ Cube(y) ∧ x ≠ y]
>
> ∃x ¬Cube(x)

Proof: Here we will use the method of **existential instantiation**. Premise 3 tells us that there are at least two cubes; let's call these *a* and *b*. Premise 1 tells us that every cube is left of something, so we can infer that if *a* is a cube, then there is something that *a* is to the left of. But *a* **is** a cube, so we may infer (using **→ Elim**) that there is something that *a* is to the left of. Let's pick an object that *a* is to the left of and call it *c*. We will now prove that *c* is not a cube. For, suppose *c* is a cube; then *a* is a cube and *c* is a cube and *a* is to the left of *c*. We may then apply ∃ **Intro** to this and infer that there is a cube that is to the left of a cube, i.e., ∃x ∃z (Cube(x) ∧ Cube(z) ∧ LeftOf(x, z)), contradicting premise 2. Since our assumption that *c* is a cube has led to a contradiction, we may infer that *c* is not a cube. So, by ∃ **Intro**, there is at least one object that is not a cube. Now we know from premise 3 that there are at least two cubes, and we have been assuming that their names are *a* and *b*. From this assumption we derived the conclusion that there is at least one non-cube. But nothing in our proof depended on the names of the two cubes—no matter what their names are, we could still derive the same conclusion. Hence we may conclude (citing premise 3 in place of our assumption about the names of the cubes) that there is at least one object that is not a cube.

§ 12.4 Proofs involving mixed quantifiers

In both of these methods (existential instantiation and universal generalization), we must be sure that we have a way of guaranteeing that the name we use in our assumption picks out an arbitrary object. We must be sure that we do not smuggle into our proof any extraneous information that may depend on the individual denoted by the name we have chosen to represent a random object.

A special case in which this issue arises involves proofs with **mixed quantifiers**. Although ∃∀ sentences imply their ∀∃ counterparts, the converse is not always true. The following arguments illustrate this point.

A valid argument

There is a certain person who is admired by everyone. Therefore, everyone admires someone or other.

> ∃y ∀x Admires(x, y)
>
> ∀x ∃y Admires(x, y)

This argument is valid, and our method of proof can establish its validity.

Proof

> There is a certain person who is admired by everyone. Let's call him *George*. Now since everyone admires George, we can pick any person at random and that person will admire George. So let's pick someone at random, and call him *Jerry*. So, since Jerry admires George, it follows that there is someone whom Jerry admires. But Jerry was any randomly chosen individual, and we have shown that Jerry admires someone or other. Therefore, everyone admires someone or other.

Switch the premise and conclusion, however, and the argument becomes invalid.

An invalid argument

Everyone admires someone or other. Therefore, there is a certain person who is admired by everyone.

> $\forall x \, \exists y \, \text{Admires}(x, y)$
>
> $\exists y \, \forall x \, \text{Admires}(x, y)$

This argument is clearly invalid. (Just to convince yourself of its invalidity, see whether you can describe a possible situation—say, involving George, Jerry, and Elaine—in which the premise is true and the conclusion false. You should be able to do this fairly easily—to check out your counterexample, compare it with mine.) So we must take pains to ensure that the method of proof we used above will **not** enable us to "prove" the conclusion of this argument. If we are not careful, that is just what will happen. Consider the following pseudo-proof:

Pseudo "proof"

> Everyone admires someone or other. So let's pick anyone at random, and call her *Elaine*. Now Elaine admires someone or other; so let's pick a person that Elaine admires and call him *Cosmo*. So, Elaine admires Cosmo. But Elaine was any randomly chosen individual, and we have shown that Elaine admires Cosmo. So everyone admires Cosmo. Therefore, it follows that there is a certain person (namely, Cosmo) whom everyone admires.

Clearly, this "proof" is fallacious. The fallacy lies in concluding, from the fact Elaine was chosen at random, and admires Cosmo, that **everyone** admires Cosmo. The reason it is a fallacy is that we did not choose Cosmo at random from the entire domain of people who are admired; rather, we chose him from among the people that **Elaine** admires. That is, our choice of Cosmo depended on our prior choice of Elaine.

This is an example of what *LPL* calls a "hidden dependency." Since the sentence *Elaine admires Cosmo* mentions an individual (Cosmo) whose choice depended on our prior choice of Elaine, we cannot universally generalize with respect to Elaine and conclude *Everyone admires Cosmo*.

As *LPL* explains(p. 331), "we must now add the restriction that $S(c)$ not contain any constant introduced by existential instantiation after the introduction of the constant c."

Here's another way of putting the restriction: even if we prove $S(c)$, where c is a randomly chosen individual, we may not draw the conclusion $\forall x \, S(x)$ if $S(c)$ contains *any* **other** name that was introduced as an assumption for existential instantiation after the name c was introduced.

Thus, in our example, S(c) is *Elaine admires Cosmo* and c is *Elaine*. Since *Elaine admires Cosmo* contains the name *Cosmo*, which was introduced after the name *Elaine was* introduced for existential instantiation strategy, we may not infer *Everyone admires Cosmo*.

This restriction *sounds* complicated, but in system $\mathcal{F}$ we will have a very simple method of ensuring that it is not violated. It will turn out that our rule of ∃ **Elim** (the formal rule that corresponds to the method of existential instantiation) will be set up in such a way as to prevent this fallacious inference. We will see *how* this comes about when we introduce the formal rules for quantifiers in Chapter 13.