Chapter 14: More on Quantification

§ 14.1 Numerical quantification

In what we've seen so far of FOL, our quantifiers are limited to the universal and the existential. This means that we can deal with English quantifiers like *everything* and *something*. We quickly discovered that with a judicious use of truth-functional connectives, we could also express such English quantifiers as *nothing*, *every cube*, *some tetrahedron*, *all large cubes*, etc. The FOL representations of these quantifier phrases in English make use of the quantifiers plus some truth-functional machinery:

Nothing	$\forall x \neg \ \ldots$
Every cube	$\forall x\ (\mathrm{Cube}(x) \rightarrow \ \ldots$
Some tetrahedron	$\exists x\ (\mathrm{Tet}(x) \land \ \ldots$
All large cubes	$\forall x\ ((\mathrm{Cube}(x) \land \mathrm{Large}(x)) \rightarrow \ \ldots$

We will now see how to use our existing FOL machinery to represent **numerical** quantifiers. At this point, there is one numerical quantifier we know how to express—*at least one*. That is because the sentence *At least one cube is large* goes easily into FOL as $\exists x\ (\mathrm{Cube}(x) \land \mathrm{Large}(x))$.

Now we will learn how to use FOL to express such numerical quantifiers as the following: *at least two*, *at most one*, *exactly one*, *at least three*, *at most two*, *exactly two*, etc. The interesting aspect of this is that we do not need to enrich FOL in any way in order to accomplish this. We will not have special "numerical" quantifiers. Rather, we will use our regular universal and existential quantifiers, together with truth-functional connectives and (most importantly) the identity sign.

In the following examples, we will be using FOL to say something about the number of cubes there are.

At least two

Suppose we want to say that there are at least two cubes. A first effort might be $\exists x \exists y\ (\mathrm{Cube}(x) \land \mathrm{Cube}(y))$. But we quickly realize that this cannot be correct. For nothing in this FOL sentence tells us that *x* and *y* have to be **different** cubes. (If this is not clear, put this sentence into a new Tarski's World sentence file and evaluate it in a world with a single cube. You will see that it comes out true. If you don't see why, try playing the game against Tarski, committing to the falsity of this sentence. Notice why Tarski will always win.)

Obviously, what is needed is a clause guaranteeing that *x* and *y* are distinct objects. And such a clause is easy to come by: $x \neq y$. So, our final version of *there are at least two cubes* is:

$$\exists x \exists y\ (\mathrm{Cube}(x) \land \mathrm{Cube}(y) \land x \neq y)$$

Notice that we can have an *at least two* quantification that is not restricted to cubes. If we want to say simply that there are at least two **things** (without being specific about any other properties these things might have) we can write:

$$\exists x \exists y\ x \neq y$$

At most one

Representing *There is at most one cube* in FOL is a bit more complicated. Here is the general idea. Suppose that your domain of discourse is a barrel that contains cubes, tetrahedra, and dodecahedra (a kind of three-dimensional Tarski's World!) and suppose that there is at most one cube in the barrel.

Now suppose that you reach into the barrel and pull out a cube, and then throw the cube back in. Then you reach in again and pull out a cube. (Pretty amazing, considering that there is at most one cube in the barrel.) In fact, we know for certain that you pulled out the same cube twice!

This is how we will put the claim when we couch it in FOL: if you reach in the barrel and pull out a cube, x, and (after returning the cube to the barrel) reach in again and pull out a cube, y, then $x = y$.

$$\forall x \forall y\ ((\mathrm{Cube}(x) \land \mathrm{Cube}(y)) \rightarrow x = y)$$

Similarly, we can have an *at most one* quantification that is not restricted to cubes. If we want to say simply that there is at most one **thing** (without being specific about any other properties it might have) we can write:

$$\forall x \forall y\ x = y$$

Exactly one

Having dealt with *at least one* and *at most one*, we already have everything we need to handle *exactly one*. For it is nothing more than the conjunction of the other two. That is, for it to be true that there to be exactly one cube is just for it to be true both that there is at least one cube and that there is at most one cube.

So a simple way to arrive at an FOL translation of *There is exactly one cube* is just to conjoin our FOL versions of *at least one* and *at most one*:

$$\exists x\ \mathrm{Cube}(x) \land \forall x \forall y\ ((\mathrm{Cube}(x) \land \mathrm{Cube}(y)) \rightarrow x = y)$$

However, there are equivalent, but more compact, ways of expressing this. We will be led to one such formulation by the following line of thought. For there to be exactly one cube is for there to be something, x, such that x is a cube, and nothing but x is a cube. That is, there is an x such that x is a cube, and no matter which y you pick, if y is a cube, then y and x are one and the same object. In FOL symbols:

$$\exists x\ (\mathrm{Cube}(x) \land \forall y\ (\mathrm{Cube}(y) \rightarrow y = x))$$

This is the version presented in *LPL* on p. 370. An even more compact version can be produced, however. We can delete the clause $\mathrm{Cube}(x)$, but get the effect of including it by changing the $\rightarrow$ to a $\leftrightarrow$. That gives us:

$$\exists x \forall y\ (\mathrm{Cube}(y) \leftrightarrow y = x)$$

In other words, to say that there is exactly one cube is to say that there is an x such that no matter which y you pick, y is a cube iff y and x are one and the same object. This is the most compact version of *exactly one*.

As before, we can have an *exactly one* quantification that is not restricted to cubes. If we want to say simply that there is exactly one **thing** (many philosophers have actually believed this!) we can write:

$$\exists x \forall y\ y = x$$

In fact, we met this sentence earlier, when we first studied multiple quantification with "mixed" quantifiers. You might wish to test this sentence out in Tarski's World. You will quickly discover that it is true in any world containing exactly one block. As soon as you add a second block, the sentence becomes false.

At least three

At least two required two quantifiers and a non-identity clause. So it is easy to see that *at least three* will require three quantifiers and three non-identity clauses. That is, in FOL we express *there are at least three cubes* as:

$$\exists x \exists y \exists z\ (\mathrm{Cube}(x) \land \mathrm{Cube}(y) \land \mathrm{Cube}(z) \land x \neq y \land y \neq z \land x \neq z).$$

Three non-identity clauses are required because we need to state that we can select cubes in such a way that after three selections, we never selected the same cube twice. And we can say simply that there are at least three things—just drop the Cube wffs from the sentence above:

$$\exists x \exists y \exists z\ (x \neq y \land y \neq z \land x \neq z).$$

At most two

To understand our treatment of *there are at most two cubes*, put yourself back in the position of someone pulling blocks out of a barrel. If there are at most two cubes in the barrel, that means that if you make three draws from the barrel (under the conditions described earlier) and get a cube every time, then you must have drawn the same cube more than once. That is, in FOL we express *there are at most two cubes* as:

$$\forall x \forall y \forall z\ ((\mathrm{Cube}(x) \land \mathrm{Cube}(y) \land \mathrm{Cube}(z)) \rightarrow (x = y \lor y = z \lor x = z)).$$

As before, we can say simply that there are at most two things by dropping the Cube wffs and the $\rightarrow$ from the sentence above and keeping just the quantifiers and the disjunction of identity clauses:

$$\forall x \forall y \forall z\ (x = y \lor y = z \lor x = z).$$

Exactly two

To handle *exactly two*, we can build on our treatment of *exactly one*. Conceptually, the simplest treatment is just to conjoin our two FOL sentences that translate *at least two* and *at most two*. But the resulting FOL sentence is not very compact. To get a more compact version, we can build on the compact version of *exactly one*. That is, instead of saying *there is something such that it and it alone is a cube*, we would say *there are two distinct things such that they and they alone are cubes.* In FOL:

$$\exists x\ \exists y\ (x \neq y \land \forall z\ (\mathrm{Cube}(z) \leftrightarrow (z = x \lor z = y))).$$

If you are a dualist, and wish to say that there are exactly two things, you'd put it this way in FOL:

$$\exists x\ \exists y\ (x \neq y \land \forall z\ (z = x \lor z = y)).$$

Examples

Let's test our understanding of some numerical sentences. Open the file Ch14Ex1.sen (on the Supplementary Exercises web page). Notice that the English sentences all make numerical claims, and that they all appear with their FOL translations. Can you see why the FOL sentences say what their English translations do?

Try constructing a world in which all of the sentences are true. Now try making changes to the world, falsifying some of the sentences. Now make more changes, so that all the sentences are false.

Next, destroy all the blocks in the world in which you are evaluating the sentences, and open the file Ch14Ex1a.sen. You will find the same FOL sentences as in the previous sentence file, but all the English translations have been deleted. Do you still know what the FOL sentences mean? Try to rebuild your world so that all the sentences come out true.

As a final test of your understanding of numerical quantification in FOL, open the file Ch14Ex2.sen. There are eight numerical claims here, but no accompanying English translations. These sentences are more complicated than the ones in the previous file, so read them carefully. Then create a world making as many of the sentences true as you can. (You should be able to create a world in which they are all true.)

You may have trouble understanding some of these sentences. Try putting any that you are in doubt about into English. To assist you with the translation, make changes to your world and observe what happens to the truth value of the sentence in question. That should help you figure out what the sentence means. Save your world as World Ch14Ex2.wld. Save the sentence file (with your English annotations) as Sentences Ch14Ex2.sen.

Generalizing

Obviously, what we have done for the numbers 1, 2, and 3, we can do for any integer *n*. That is, for any *n*, we can produce FOL sentences that translate:

There are at least *n F*'s

There are at most *n F*'s

There are exactly *n F*'s.

Needless to say, as *n* gets larger, the FOL sentences become longer and more complex—this is hardly an ideal language in which to do arithmetic! But the point is that the expressive power of FOL is considerable. For any condition expressible in FOL and for any finite number, *n*, one can, in principle, construct an FOL sentence saying that *n* things satisfy that condition.

Abbreviations for numerical claims

Rather than write out the (sometimes very long) FOL sentences that express numerical claims, we can use the following **abbreviation** scheme.

- $\exists^{\geq n}$x P(x) abbreviates the FOL sentence asserting "There are at least *n* objects satisfying P(x)."
- $\exists^{\leq n}$x P(x) abbreviates the FOL sentence asserting "There are at most *n* objects satisfying P(x)."
- $\exists^{!n}$x P(x) abbreviates the FOL sentence asserting "There are exactly *n* objects satisfying P(x)."

For the special case where $n = 1$, it is customary to write ∃!x P(x) as a shorthand for $\exists^{!1}$x P(x). This can be read as "there is a unique *x* such that *P*(*x*)."

More translations

In translating numerical claims, we made heavy use of =, the identity predicate. There are other common claims, not explicitly numerical, that also require the use of identity. We will look at a couple of them here.

Superlatives

A superlative is an adjective ending in *–est*, such as *largest*, *oldest*, *strongest*, etc. Suppose we want to write an FOL sentence corresponding to *b is the largest cube*. We might try one of the following:

1. $\text{LargestCube}(b)$
2. $\text{Cube}(b) \wedge \text{Largest}(b)$

But both of these seem problematic. (1) conceals too much information, for it does not have $\text{Cube}(b)$ as a FO consequence, whereas *b is a cube* certainly seems like a FO consequence of *b is the largest cube*. (2) avoids this problem, but introduces another. For (2) says that *b* is the largest *thing*, whereas our original sentence only says that *b* is the largest *cube*.

And *b* might be the largest cube without being the largest thing. (Imagine a world in which *b* is a medium cube, all the other cubes are small, and *c* is a large tetrahedron.)

The trick is to translate this sentence into FOL using only the *comparative* predicate *larger*. To say that *b* is the largest cube is to say that *b* is larger than all the *other* cubes. That is, *b* is a cube, and every cube that is not *b* is smaller than *b*. In FOL:

$$\text{Cube}(b) \wedge \forall x\, ((\text{Cube}(x) \wedge x \neq b) \rightarrow \text{Larger}(b, x))$$

In general, to be the *F*-est thing is to be *F*-er than everything else; to be the *F*-est *G* is to be a *G* that is *F*-er than every **other** *G*. In colloquial speech, people tend to be careless and leave out the *other*. Meaning to assert that Clark Kent is the strongest man, they may say *Clark Kent is stronger than any man*. Strictly speaking, of course, this is not true: Clark Kent may be stronger than all the *other* men, but he is not stronger than *himself*!

Exceptives

An *exceptive* is a claim that makes a universal generalization with an exception, such as *everything is a cube except c*. We can translate *everything is a cube* into FOL as $\forall x\, \text{Cube}(x)$, and *c is not a cube* as $\neg\text{Cube}(c)$. But if we simply conjoin these:

$$\forall x\, \text{Cube}(x) \wedge \neg\text{Cube}(c)$$

we get a contradiction, which our original sentence certainly is not. What we want to say, roughly, is *c is not a cube, but everything* ***else*** *is a cube*:

$$\neg\text{Cube}(c) \wedge \forall x\, (x \neq c \rightarrow \text{Cube}(x)).$$

This translation is correct, but we can produce a more compact version. Think of the sentence *c is not a cube* as a universal generalization: *c is no cube*, or (equivalently) *no cube is c*. So instead of $\neg\text{Cube}(c)$ we can write:

$$\forall x\, (\text{Cube}(x) \rightarrow x \neq c).$$

Using this in place of the equivalent $\neg\text{Cube}(c)$, we get:

$$\forall x\, (\text{Cube}(x) \rightarrow x \neq c) \wedge \forall x\, (x \neq c \rightarrow \text{Cube}(x))$$

which, by moving the quantifier to the outside (see §10.4), is equivalent to:

$$\forall x\, [(\text{Cube}(x) \rightarrow x \neq c) \wedge (x \neq c \rightarrow \text{Cube}(x))].$$

Finally, we replace the embedded conjunction of conditional wffs with the corresponding biconditional, and obtain:

$$\forall x\ (\text{Cube}(x) \leftrightarrow x \neq c)).$$

This is the simplest way to express the exceptive sentence *everything is a cube except c* in FOL. Now read this FOL sentence from left to right, and compare it with the English sentence. Did you notice that the phrase *except c* is rendered in FOL by $\leftrightarrow x \neq c$? Exceptive sentences go into FOL as *negative biconditionals*.

§ 14.2 Proving numerical claims

Since we can translate numerical claims into FOL, we can evaluate the validity, in FOL, of arguments containing such claims. Consider, for example, the following argument

| There are exactly two cubes.
| There are exactly three non-cubes.
|—
| There are exactly five objects.

The conclusion is clearly a logical consequence of the premises—there is no possible circumstance in which there are two cubes, three non-cubes, but not a total of five objects, cubes and non-cubes, altogether. But is the conclusion a **FO consequence** of the premises?

Using our abbreviations for FOL numerical quantifications, the FOL version of the argument looks like this:

| $\exists^{!2}x\ \text{Cube}(x)$
| $\exists^{!3}x\ \neg\text{Cube}(x)$
|—
| $\exists^{!5}x\ (x = x)$

Can we prove this conclusion in $\mathcal{F}$? Before we can do so, of course, we would have to write out the "real" FOL sentences, instead of the abbreviations we used above. And when we do this, we will see that the argument contains no predicates that affect its validity—Cube could just as easily be replaced by Tove or any other predicate. So the conclusion is, indeed, a FO consequence of the premises. And since $\mathcal{F}$ is complete, it follows that it is possible, at least in principle, to prove this conclusion in $\mathcal{F}$.

What this means

Such a proof would seem to come very close to being a proof, **purely within formal logic**, that 2+3=5. Of course, we are not really proving things about *numbers*, but about cubes and non-cubes, and about relations among various conditions of their identity and distinctness. But we are certainly capturing some basic arithmetical ideas within a system of pure logic.

What this does not mean

Although we can express (and prove) many arithmetical claims in FOL, there are still many kinds of arithmetical claims that we can neither express nor prove. For example, consider the following claim:

There are more cubes than tetrahedra.

This claim does not tell us how many cubes there are—only that the number of cubes is larger than the number of tetrahedra. In other words, there are numbers *n* and *m* such that there are *n* cubes and there are *m* tetrahedra, and $n > m$. One might try to state this in FOL as follows:

$$\exists n \exists m\, (\exists^{!n} x\ \mathrm{Cube}(x) \wedge \exists^{!m} x\ \mathrm{Tet}(x) \wedge n > m)$$

But there are two problems here. First, our sentence contains the predicate >, which is not one of the logical symbols of FOL. Second (and more importantly), what we have written is not the abbreviation of any FOL sentence. A numerical subscript of the form $^{!n}$ tells us how to write the FOL sentence we are abbreviating, but only when n is some positive integer. We have no way of dealing with a **variable** in a numerical quantifier.

In short, our numerical quantifiers do not quantify over numbers; they are simply abbreviations of more complex looking FOL sentences that quantify over whatever objects (cubes, etc.) are in their domain of discourse. We can **axiomatize** arithmetic in FOL (see. §16.4), but we cannot express all arithmetic claims in "pure" FOL.

In practice, of course, proofs of arithmetical claims in FOL are messy and difficult to come by—FOL is not a very good language in which to do arithmetic. Still, we can easily handle cases where n is very small, and use Fitch to construct proofs of at least some numerical claims. Here's an easy example to start with:

| There is at least one cube.
| There is at least one tetrahedron.
|—
| There are at least two things.

You'll find the problem on the Supplementary Exercises page as Ch14Ex3.prf. In this problem, you'll need to use **Ana Con**. You should use it only to obtain ⊥ from atomic sentences.

§ 14.3 *The*, *both*, and *neither*

When the word *the* combines with a noun phrase to form an expression that purports to refer to exactly one object, the entire phrase (of the form *the so-and-so*) is called a **definite description**. Here are some examples of definite descriptions:

The tallest player on the team

The king of Norway

The sum of 3 and 5 [more usually written '3 + 5']

The 40th president of the U.S.A.

Whistler's mother [note the absence of *the* in this case]

Notice that definite descriptions function **syntactically** like names, as illustrated by the following pairs of sentences:

John has red hair.

The tallest player on the team has red hair.

Reagan was a Republican.

The 40th president of the U.S.A. was a Republican.

But there is good reason to think that definite descriptions do not function **semantically** like names. In fact, FOL would be inadequate if it treated descriptions in the same way it treats proper names, namely, as logical constants. For then the FOL versions of both of these arguments would be, effectively, the same:

| John has red hair.
|—
| Some player on the team has red hair.

The tallest player on the team has red hair.

Some player on the team has red hair.

Clearly, the second argument is valid, but the first is not (for nothing in the first argument tells us that John is a player on the team). The premise of the second argument contains information that the premise of the first argument lacks. So we should not treat definite descriptions in FOL as if they were names.

But now a problem arises. For logic cannot guarantee that a definite description actually succeeds in picking out a unique object. Consider a sentence like *The cube is small,* and imagine that you are evaluating it in various Tarski Worlds. How would you assess its truth value in a given world? You would expect to find exactly one cube in the world, and then you would check its size—if it's small, the sentence is true; otherwise, the sentence is false.

But suppose there are no cubes in the world? What is the truth value of the sentence in that case? Or suppose there are two cubes, one of which is small and one of which is not? What is the truth value of the sentence in that case?

In both of these cases, something has gone wrong with the description. For convenience, let's say that a description, *the F*, is a **good** description when there is exactly one *F*, and a **bad** description otherwise. Thus, there are two ways in which a description can go bad. *The senator from Washington* is a bad description in one way, since there is more than one senator from Washington, and *the present king of France* is a bad description in another way, since there is no king of France at present.

How are we to evaluate sentences that contain bad descriptions? This was the problem that motivated Bertrand Russell's famous Theory of Descriptions (1905).

Russell's Theory of Descriptions

According to Russell, a sentence containing a definite description can be thought of as a conjunction with three conjuncts. Consider such a sentence:

> The cube is small.

On Russell's theory, this amounts to the following conjunction:

> There is at least one cube, and there is at most one cube, and every cube is small.

This easily goes into FOL as:

$$\exists x\ \mathrm{Cube}(x) \land$$
$$\forall x \forall y\ ((\mathrm{Cube}(x) \land \mathrm{Cube}(y)) \rightarrow y = x) \land$$
$$\forall x\ (\mathrm{Cube}(x) \rightarrow \mathrm{Small}(x))$$

It's easy to see that this sentence can be false in three different ways, depending on which conjunct is false: there may be no cubes (first conjunct false), or more than one cube (second conjunct false), or some cube that is not small (third conjunct false).

Equivalent formulations of Russell's analysis

Equivalently, but more compactly, Russell's analysis can put as follows:

$$\exists x\ (\mathrm{Cube}(x) \land \forall y\ (\mathrm{Cube}(y) \rightarrow y = x) \land \mathrm{Small}(x))$$

This is the “standard” FOL sentence that *LPL* presents as Russell’s analysis of *the cube is small*.

An even more compact version looks like this:

$$\exists x\ \forall y\ ((\text{Cube}(y) \leftrightarrow y = x) \wedge \text{Small}(x))$$

That is, there is exactly one cube, and it’s small. All three versions are, of course, equivalent.

Both **and** ***neither***

We may note in passing that Russell’s analysis can be extended to cover the determiners *both* and *neither*. That is, we can treat phrases like *both cubes* and *neither tetrahedron* along the same lines as *the cube*. This can be seen easily from the following examples.

Both cubes are small.

On Russell’s analysis, this says that there are exactly two cubes, and each cube is small. That is:

$$\exists^{!2} x\ \text{Cube}(x) \wedge \forall x\ (\text{Cube}(x) \rightarrow \text{Small}(x))$$

Similarly, *neither tetrahedron is large*, on Russell’s analysis, says that there are exactly two tetrahedra, and no tetrahedron is large. That is:

$$\exists^{!2} x\ \text{Tet}(x) \wedge \forall x\ (\text{Tet}(x) \rightarrow \neg\text{Large}(x))$$

Remember that what we have produced above are really just *abbreviations* of the real FOL sentences that would count as the Russellian analyses of *both* and *neither*. (Real FOL sentences don’t contain numerical quantifiers, like $\exists^{!2}x$.)

Two key features of Russell’s theory

The beauty of Russell’s analysis is that it provides a truth value for every sentence containing a definite description, even if it’s a bad description. If someone says *the cube is small* when there is no cube, he or she has simply said something false.

Russell’s analysis also provides this interesting feature: although a sentence containing a description may be perfectly unambiguous, the introduction of a logical operation such as negation may introduce an **ambiguity**. Thus, consider the sentence *the cube is small*, whose Russellian analysis looks like this:

1. $\exists x\ (\text{Cube}(x) \wedge \forall y\ (\text{Cube}(y) \rightarrow y = x) \wedge \text{Small}(x))$.

Now consider what happens when a negation is introduced:

The cube is not small.

On Russell’s theory, this sentence is ambiguous. On one reading, it asserts that there is exactly one cube, and says, further, that it is not small. In FOL:

2. $\exists x\ (\text{Cube}(x) \wedge \forall y\ (\text{Cube}(y) \rightarrow y = x) \wedge \neg\text{Small}(x))$

But on another reading, the English sentence says something different, namely, that it is **not** the case both that there is exactly one cube and that it is small. In FOL:

3. $\neg\exists x\ (\text{Cube}(x) \wedge \forall y\ (\text{Cube}(y) \rightarrow y = x) \wedge \text{Small}(x))$

You can see the difference between these sentences by comparing their evaluations in various worlds. Open the file Russell.sen, where you will find these three sentences. Now create a world with a single cube in it (it may contain non-cubes, too, but they are irrelevant to these sentences) and evaluate all three sentences.

You will notice that (2) and (3) will always agree with one another, and disagree with (1), so long as the description is "good"—i.e., as long as there is exactly one cube in the world. But notice what happens when you add a cube, or remove all the cubes. In worlds like this, the description *the cube* is "bad," and sentences (2) and (3) will diverge in truth value. Sentence (1) will be false; but **(2) will also be false**. (After all, (2) is not the negation of (1), since (2) has its ¬ embedded in the last conjunct.) The negation of (1) is (3), and **(3) will be true**. On Russell's theory, then, there are situations in which both *the cube is small* and *the cube is not small* are false.

Strawson's analysis

Russell's theory has not convinced everyone. One celebrated critique is that of philosopher P. F. Strawson. According to Strawson, Russell is mistaken in supposing that one who utters the sentence *the cube is small* makes three claims—that there is at least one cube, and at most one cube, and that every cube is small. Rather, such a person does not even succeed in making a claim unless there is exactly one cube. That there is exactly one cube is not part of what the speaker claims, but is a **presupposition** of his making a claim at all.

If the presupposition is fulfilled (that is, if there is exactly one cube), then the utterer of the sentence *the cube is small* claims, about that cube, that it is small. If the presupposition is **not** fulfilled (that is, if there is more than one cube, or if there are no cubes), then the speaker has failed to make any claim.

So one who utters a sentence containing a bad description has not succeeded in making a claim. And on Strawson's account, truth values attach not to the sentences we utter, but to the claims we make with them. Hence, according to Strawson, nothing true or false gets expressed by a sentence containing a bad description. Strawson's analysis thus introduces what have been called **truth value gaps**.

Consequences of Strawson's analysis

An obvious consequence of Strawson's analysis is that sentences like *the cube is small* cannot be translated into FOL. For all FOL sentences have truth values, at least if they do not contain any names. Strawson's proposal is, in effect, to treat definite descriptions (semantically as well as syntactically) in the way that FOL treats names: the sentences in which they occur cannot be evaluated for truth value unless they succeed in uniquely referring.

Strawson's analysis of definite descriptions seriously weakens FOL. For FOL would be unable, on Strawson's account, to explain the validity of certain obviously valid arguments. Consider this example:

| The large cube is in front of *b*.
|—
| Something large is in front of *b*.

Although Strawson would agree that this is a valid argument (there is no possible circumstance in which the premise is true and the conclusion false), the conclusion would not be a FO consequence of the premise in a Strawsonian FOL. For it would look like this:

$$
\begin{array}{|l}
\text{FrontOf(the large cube, b)} \\
\hline
\exists x\,(\text{Large}(x) \land \text{FrontOf}(x, b))
\end{array}
$$

And the conclusion of this argument is not a FO consequence of the premise. Russell's theory, of course, proposes a different translation into FOL:

$$
\begin{array}{|l}
\exists x \forall y(((\text{Large}(y) \land \text{Cube}(y)) \leftrightarrow y = x) \land \text{FrontOf}(x, b)) \\
\hline
\exists x\,(\text{Large}(x) \land \text{FrontOf}(x, b))
\end{array}
$$

And the conclusion here is obviously a FO consequence of the premise. Can Russell's theory—which does not allow Strawson's "truth value gaps"—be defended?

Response to Strawson

It is Strawson's notion of *presupposition* that introduces truth value gaps. So to do away with them, we need an alternative account of what he calls presuppositions. The best alternative is to say that these are really **implicatures**.

According to Strawson, the sentence *the large cube is not in front of* ***b*** carries with it the presupposition that there is exactly one large cube. When that presupposition is not fulfilled, one who utters the sentence fails to make any claim, true or false. But if it is only an implicature that there is exactly one large cube, the sentence may still have a truth value even when the implicature is false.

So let us apply the **cancellability** test. Can one conjoin *there is not exactly one large cube* to *the large cube is not in front of* ***b*** without contradiction? Opinions are mixed. My view is that there is no contradiction here. It is not self-contradictory to say "The large cube is not in front of *b*—in fact, there is no large cube at all!"

It is rather like saying to the child who thinks that the tooth fairy put a dollar under her pillow: "No, the tooth fairy did not put a dollar under your pillow—in fact, there *is* no tooth fairy." It may be unkind to say this, but it is not untrue. Obviously, then, it is not self-contradictory, either.

In fact, a Strawsonian analysis of this case seems particularly wrong-headed. Consider his account of the sentence uttered by the child:

The tooth fairy put a dollar under my pillow.

According to Strawson's account, this sentence cannot be used to make a claim unless its presuppositions are fulfilled. And one of its presuppositions is that there is exactly one tooth fairy. So, since there are no tooth fairies, the child has not made a claim at all.

But surely this is wrong. The child *has* made a claim, and a false one. And we can correct the child (if we're mean enough, or the child is old enough) by uttering the proper negation of what the child said:

No, the tooth fairy did not put a dollar under your pillow.

This is the harsh truth, and the Russellian analysis gives just the right result.