What is the Frege/Russell Analysis of Quantification?
Scott Soames

The Frege-Russell analysis of quantification was a fundamental advance in semantics

and philosophical logic. Abstracting away from details idiosyncratic to each, the key idea is

that quantifiers express higher level properties that are predicated of lower level properties.

The two philosophers express this *properties of properties* analysis slightly differently.

Whereas Russell talks of properties of propositional functions (which are functions from

objects to structured propositions), Frege speaks of senses that present higher-level concepts

that are predicated of lower-level concepts. Since the issues I will raise are neutral between

these different frameworks, I will stick with property talk.[1]

The analysis is illustrated by claims P1 – C3 about sentence 1.

1. ∀x Fx / Everything is F

P1. The proposition expressed by (1) consists of two constituents – the property *being true of everything* and the property *being F* – the first of which is predicated of the second.

P2. If a proposition p consists of two constituents – *being so-and-so* and *such-and-such*, the first of which is predicated of the second, then p = the proposition *that such and such is so-and-so.*

C1. The proposition expressed by (1) = the proposition *that being F is true of everything.* (Call this proposition 'Prop 1'.)

P3. Prop 1 is also expressed by sentence (1a).

1a. The property *being F* is true of everything.

C2. So, (1) and (1a) express the same proposition.

The first thing to notice about my illustration is that although I have called the account

of quantification "an analysis," the sentences used to specify the proposition expressed by

---

[1] 'F' is used as a schematic letter throughout both this argument and the rest of the text.

sentence (1) itself contains the universal quantifier being analyzed. That may seem strange.

Indeed, it may seem worse than strange, as the following argument shows.

P4.   The expression 'is true of' that occurs in sentence (1a) is a 2-place predicate that expresses the relation *being true of*, which holds between properties and objects. Since the quantifier 'everything' stands in the second argument-place of this predicate, (1a) is a universally quantified sentence that expresses the same proposition as (1∀).

1∀.   Everything is such that *the property being F is true of it.* (∀x: *the property being F is true of* x.)

P5.   So, the Frege/Russell analysis of quantification applies to (1a/1∀), yielding a proposition that has two constituents the first of which – *being true of everything* – is predicated of the second – *being an object of which the property F is true.*

P6.   This proposition (described in P5), is the proposition *that the property being an object of which the property being F is true* is true of everything. Call it 'Prop 2'.

C3.   The proposition expressed by (1a/1∀) is Prop 2.

C4.   Since (1) and (1a) each express just one proposition, Prop 1 = Prop 2.

The cycle can be repeated producing a hierarchy of propositions each of which is identified with all the others.

Prop 1 The proposition that *being F* is true of everything.

Prop 2 The proposition that *being an object of which (being F) is true* is true of everything.

Prop 3 The proposition that *being an object of which (being an object of which being F is true) is true* is true of everything.

.
.
.

    Is this result problematic? The propositions in the sequence are necessarily and apriori equivalent to one another. There is even some temptation to think that anyone who asserts or believes Prop 1 should be counted as also asserting or believing Prop 2. After all, it might be argued, anyone who *explicitly* accepted, and thereby believed, the former would be *disposed* to accept, and rightly be counted as believing, the latter. Even so, however, we don't get the needed result – namely, that the propositions in the hierarchy are *identical* with one another. We don't even establish that any anyone who believes or asserts one of them

believes or asserts all the others. Although there is some justification for thinking that one who *explicitly* asserts or judges, and hence *occurrently* believes, p may be counted as *implicitly* asserting and believing *obvious* necessary and apriori consequences of p, this doesn't justify taking assertion and belief to be closed under obvious necessary and apriori consequence. Of course, if it is in the nature of propositions that those that are necessarily and apriori equivalent are *identical*, then the hierarchy I have generated is simply a harmless repetition of the same proposition over and over again. However, since Frege and Russell were proponents of structured propositions, this wouldn't have impressed them.

Those like me who share their belief in structured propositions face a more serious problem. The constituents of Prop 1 are the properties *being true of everything* and *being F*. The constituents of Prop 2 are *being true of everything,* and *being an object of which the property being F is true.* Note, the property *being F* is itself a constituent of the property *being an object of which the property being F is true,* and so is a sub constituent of Prop 2. However, *being F* is not the predication target in Prop 2, as it is in Prop 1. In Prop 2, the target, of which *being true of everything* is predicated, is *being an object of which the property being F is true.* Since the predication target is *being an x such that the property being F is true of x,* identifying Props 1 and 2 would require identifying the simple property *being F* with a relational property of which *it* is merely a constituent. Since this is absurd, proponents of structured propositions must distinguish the two propositions, along with all others in the hierarchy. But this is intolerable; the unique proposition expressed (1) can't be identical with infinitely many different propositions.

Consequently, the proponent of structured propositions can accept the Frege/Russell analysis only by rejecting the conclusion C4 that identifies Props 1 and 2. This, in turn, requires blocking the argument before it reaches C3 (which wrongly identifies them). How

did we land in this fix? We started off by accepting an analysis of quantification that identifies the proposition p expressed by an arbitrary sentence 'Everything is F' with one that predicates *being true of everything* of the property *F*. However, since the analysis *uses* the very quantifier being analyzed, we were tempted to apply the analysis again to the sentences used to identify p. When we did this, we arrived at a new proposition p* different from p. This is intolerable if – as we have been assuming – sentence (1) expresses a single proposition. The challenge for the proponent of structured propositions is to explain why the Frege/ Russell analysis of quantification can't be applied to its own output.

The natural thought is that we went wrong in using the quantifier 'everything' in the analysis of sentences containing it. When Frege and Russell analyzed numerical concepts they were careful not to employ those concepts, but rather to define them in terms of what they took to be purely logical concepts. The same is true of Russell's analysis of descriptions. If his analysis of definite descriptions were correct, then someone who had mastered universal quantification, negation, and identity could acquire the definite article by being given its analysis. Russell seems to have understood his analysis of indefinite descriptions in the same way. His stock examples involve substitution of an adjective functioning as a predicate for the nominal in the description. Thus, 'I met a man' is always said to express the proposition *that for some x, I met x and x is human*, not *for some x, I met x and x is a man.* This is noteworthy, since for many indefinite descriptions – like 'a tiger' and 'a large and rowdy man' there is no corresponding adjectival form. There is, presumably, a reason he avoids examples like these, in which there is no obvious way of stating the analysis without using the expression being analyzed. They don't fit the conception of *analysis* he had in mind.

This point is reinforced by his analysis of quantification, summarized by R.

4

R   A sentence S_E containing 'everything' in a position capable of being occupied by a singular-term expresses the proposition *that fs is always true* – where fs is the propositional function (from objects to propositions) designated by the formula that results from replacing 'everyone' in S_E with a free occurrence of the variable 'x'. A sentence $S_∃$ containing 'something' in a comparable position expresses a the proposition *that it is not the case that fs is always false*.

The key point about Russell's formulation, which was surely no accident, is that it *does not use* the universal quantifier 'everything' in specifying the constituent that quantifier contributes to propositions expressed by sentences containing it. Instead of employing the predicate 'is true of everything', he employs the adverbially modified predicate 'is always true*'*, which is short for the more elaborate, but still putatively primitive predicate 'always assigns a true proposition to an argument'. This predicate designates a genuine property of propositional functions that is a constituent of propositions expressed by universally quantified sentences. What then does the predicate 'is true of everything' – which is short for 'assigns a true proposition to every argument' – designate? The answer, I imagine, is that it is a Russellian incomplete symbol, and so doesn't designate anything. If so, then sentence (1) will express Prop 1_R, while sentences (1a/1∀) will express Prop 2_R.

Prop 1_R: the proposition that *being F* is always true

Prop 2_R: the proposition that *being an object of which the property being F is true* is always true

(Here I revert to property rather than propositional function talk.) Although these propositions are different, no paradox results because the second is not generated from the analysis that produced the first. One can, of course, still generate a sequence of necessarily and apriori equivalent propositions, but the basis for identifying them has been eliminated.

        Russell extends his analysis to (2) – 'Something is F' – by defining 'something' in terms of the universal quantifier plus falsity and negation. This is fine when giving a system of logic. Since any choice of standard primitives plus the usual definitions yields the same

truth-preserving relations between sentences, one can adopt Russell's suggestions, or make different choices as one wishes. But since logical equivalence doesn't track identity of propositions expressed, and since there are often facts to be captured about which of the different but equivalent propositions an agent asserts or believes, semantic analyses of quantified propositions require more. For these purposes, we are better off identifying the proposition *that something is F* with the proposition that *being F is sometimes true*.

But this is only the beginning. The Russellian analysis must be extended to all generalized quantifiers, including those in (3) and (4).

3a. Some F is G.
 b. Every F is G.
 c. No F is G.

4a. Most Fs are Gs.
 b. Exactly two Fs are Gs.
 c. At least seven Fs are Gs.
 d. Many Fs are Gs.
 e. Few Fs are Gs.

The properties predicated of the propositional function (or property) *G* by the propositions expressed by the sentences in (3) are: *being sometimes true of Fs, being always true of Fs,* and *never being true of Fs. Reductio by hierarchy* is avoided as before.

> *Reductio*-Avoiding Explanation for the Examples in (3)
> The predicates 'is sometimes/always/never true of Fs*'* stand for genuine properties that are predicated of lower-level properties in propositions expressed by quantified sentences. The predicates 'is true of some/every/no Fs', designate nothing, but rather are incomplete symbols to be eliminated by analysis. Though necessary and apriori equivalents, the propositions expressed by (5) and (6) are different, with the former being primary and grasped independently of the latter.
> 5. So and so is *sometimes/always/never true of Fs.*
> 6. So and so *is true of some/every/no F.*

The quantifiers in (4) are treated similarly. The properties predicated of the propositional function (or property G) in (4a) and (4b) are *being most times true of Fs* and *being exactly two times true of Fs.* In the case of (4c) – (4e) the quantificational properties

are *being true at least seven times of Fs, being many times true of Fs,* and *being few times true of Fs.* The *reductio*-avoiding explanation is the same as before.

Although elaborating the Frege/Russell analysis of quantification in this way avoids the original *reductio,* I think it is a sham. Imagine being told, without benefit of the argument here, that "*Being G is a property that is few times true of Fs*," or that "*The propositional function that assigns to any object o the proposition that o is G is few times truth-assigning for Fs*." Far from being transparent, these remarks border on the bizarre. The natural reaction is to assimilate them to more familiar quantified forms. "Oh, I see," one is inclined to think, "You're saying that being G is a property that is true of *few Fs*," or "You are saying that the propositional function that assigns to any object o the proposition that o is G assigns a true proposition to *few Fs* (or to *few things that are F)*." The worry is not that the Russellian can't admit the speaker said this. The worry is that his explanation gets things backwards. According to him, the speaker *implicitly* said that being G is a property that is true of few F's – or that the propositional function that assigns to any object o the proposition that o is G assigns a true proposition to few things that are Fs – *by virtue of having explicitly asserted* the allegedly different, adverbially-expressed, proposition, of which the proposition implicitly asserted is a trivial consequence. Not so, one is inclined to object; it is our understanding of the Russellian adverbial formulations that is *parasitic* on our prior understanding of the quantificational locutions, not the other way around.

Although I take this objection to be correct, it is stronger than needed. The adverbial strategy for avoiding the *reductio* requires the notions *being every time true of F's, being at least one time true of F's,* and *being few (or most) times true of F's* to be primitive, conceptually prior to, and used in the analysis of, *being true of every F, being true of at least one F,* and *being true of few (or most) F's.* Surely this is misguided. If one set of notions is

7

to be explained in terms of the other, the objector is right in taking the adverbial formulations to be derivative rather than primary. But we need not insist on this. It is enough to note that *if* the adverbially expressed notions can be taken to be primitive, then those expressed using the ordinary quantifiers 'every F', 'at least one F', and 'few Fs' can too.

With this I put aside the adverbial strategy for blocking the *reductio*, and return to the practice of using quantifiers undergoing analysis to specify the propositions expressed by sentences containing them.   This involves accepting the original argument up to C2.

P1.   The proposition expressed by (1) consists of two constituents – the property *being true of everything* and the property *being F* – the first of which is predicated of the second.

P2.   If a proposition p consists of two constituents – *being so-and-so* and *such-and-such*, the first of which is predicated of the latter, then p = the proposition *that such and such is so-and-so.*

C1.   The proposition expressed by (1) = the proposition *that being F is true of everything.* (Call this proposition 'Prop 1'.)

P3.   Prop 1 is also expressed by sentence (1a).

1a.   The property *being F* is true of everything.

C2.   So, (1) and (1a) express the same proposition.

At this point, we need to find a new way of challenging P4, so as to block reductio-generating conclusion C3.

P4.   The expression 'is true of' that occurs in sentence (1a) is a 2-place predicate that expresses the relation *being true of,* which holds between properties and objects. Since the quantifier 'everything' stands in the second argument-place of this predicate, (1a) is a universally quantified sentence that expresses the same proposition as (1∀).

1∀.   Everything is such that *the property being F* is true of it. (∀x: *the property being F* is true of x.)

P5.   So, the Frege/Russell analysis of quantification applies to (1a/1∀), yielding a proposition that has two constituents the first of which – *being true of everything* – is predicated of the second – *being an object of which the property F is true.*

P6.   This proposition (described in P5), is the proposition *that the property being an object of which the property being F is true* is true of everything. Call it 'Prop 2'.

8

C3.   The proposition expressed by (1a/1∀) is Prop 2.

To block this reasoning, we must prevent the analysis from reapplying to (1a) in the manner indicated by P4 and P5. We must either allow the analysis to apply in a *different* way that does *not* produce the original hierarchy, or explain why applying it again would be a mistake.

Let's try the former. We ask, "What is the property *being true of everything*?" Suppose it is *being a property P such that for every x, P is true of x – i.e. λP (for every x: P is true of x).* So understood, the proposition expressed by (1a) is *not* identical to that expressed by (1∀) – as claimed by P4. The two are merely equivalent. On this new view, the predicate 'is true of everything' is *not* an incomplete symbol; it expresses a genuine property that is predicated of *being F*. This is an improvement, which is squarely in line with the fundamental idea behind the Frege/Russell analysis, was *not* captured by taking (1∀) to express the proposition expressed by sentences (1) and (1a). On the present suggestion, these two sentences express the proposition given by (1aλ).

1aλ.   λP [Every x: P is true of x] the property *being F*
       Being F is a property that is true of everything.

With this change, the previous hierarchy is no longer generated.

However, this doesn't help. Since (1aλ) contains 'every x' we simply get a new and equally worrisome hierarchy. We start with (1), which we are told expresses the same proposition as (1a), which, by the route just taken, expresses the same proposition as (1aλ).

1.   Every x: Fx.
1a.   The property *being F* is true of everything.
1aλ   λP [Every x: P is true of x] the property *being F*

This is trouble. Because (1aλ) contains a quantified clause, it is structurally and computationally more complex than (1a), which is itself reason to doubt that the two

9

sentences express the same proposition. The case is strengthened by applying the analysis

again inside the clause, which results in (1aλλ).

1aλλ λP [λQ (Every x: Q is true of x) *being an object of which P is true*] the property
being F

The property *being F is a property P such that the property being an object of which
P is true* is true of everything

This proposition is even more complex, containing new and more complex constituents than

proposition (1a)/(1aλ). Since it isn't identical with that proposition, we still have a *reductio*.

The problem, I suggest, is that we have misconstrued the Frege/Russell analysis of

quantification from the outset. The point of the analysis is *not* to *define* quantificational

locutions like 'everything', 'at least seven Fs', and 'Most Fs' in other terms. As we saw in

pursuing the adverbial strategy, the proposed substitutes are no advance on the originals.

Their only virtue is that because they don't contain the quantifiers under analysis, there is no

temptation to reapply the analysis, and so generate a hierarchy. But surely, taking *them* to

be primitive is no better than taking *being true of everything, being true of at least seven Fs,*

being true of most Fs to be primitive.

The correct way of understanding the Frege/Russell analysis is to see it as a theory of

logical form. Instead of specifying the propositions expressed by *every sentence* containing

quantificational locutions, it targets a *subclass* of these sentences to be explained using

quantifier-containing sentences *outside the class*. The sentences outside the class are the

analysis-giving ones. In these sentences, a quantifier Q occurs in a predicate ⌈is true of Q⌉

the argument of which is a term designating a property or propositional function applying to

the entities over which Q ranges. Sentences of this form are *antecedently intelligible without*

further analysis, and express propositions in which the property expressed by the predicate

containing the quantifier is predicated of the property or propositional function designated

by its argument. The point of the analysis is *not* to analyze the *contents* of quantifiers, but to indicate the semantic role quantifiers play as constituents of higher-level predicates.

This is a theory of *logical form* the goal of which is to reveal the structure of the propositions expressed by sentences containing quantifiers. The analysis tells us that even when a quantifier appears in a sentence in which no higher-level predication is explicit, the true *logical form* of the sentence is one in which the quantifier is used to predicate a higher-level property of the property, or propositional function, designated by the formula that results from extracting the quantifier from the original sentence. When this is understood, the explanation of why the analysis doesn't reapply to the logical form of the original sentence is evident. It doesn't apply because, once we have arrived at the logical form of a sentence, there is nothing further to do.