

C. Daly, An Introduction to Philosophical  
Methods, Buffalo 2010, Baccarini Press

## The Paradox of Analysis

model of analysis. The model in question is relatively simple. It draws upon Moore's views on analysis. Moore's views on analysis — or what have been taken to be his views — have been highly influential in analytic philosophy. Many philosophers have endorsed the model of analysis in question. Many others though have rejected it in part or as a whole. By setting out the model, we can locate various philosophical views on analysis depending upon which parts of the model they accept, which parts they reject, and what they replace them with.

The model will be set out as a model of the analysis of terms. An equivalent way of setting out the model would be as a model of the analysis of concepts.

The model makes the following five claims:

(1) *An analysis has the logical form of a universally quantified biconditional:*

$$(\forall x) (Fx \leftrightarrow Gx)$$

The above formula says that, for any thing  $x$ ,  $x$  is  $F$  if and only if  $x$  is  $G$ . As a matter of convention, the analysandum is taken to be the formula on the left-hand side of the biconditional, and the analysans is taken to be the formula on its right-hand side. An analysis gives necessary and sufficient conditions for something to be  $F$ .

(2) *An analysis is necessarily true.* The analysis does not apply only to every actual thing that is  $F$ . It applies to everything that could be an  $F$ . An analysis of, for example, the term "cause" does not apply only to every actual cause. The analysis applies to everything that could be a cause, to every possible cause. So an analysis is to be read as saying that, for every possible thing, that thing is  $F$  if and only if it is  $G$ . So, if the analysis is true, it will be necessarily true. Furthermore, if the analysis is true, it is true because of the meanings of the terms in it. Hence if the analysis is true, it is analytically true — true in virtue of the meaning of the terms involved. Let's spell this out a little.

Suppose that  $F$  and  $G$  are concepts or word-meanings, and that it is a conceptual truth that  $F$ s are  $G$ s. A conceptual truth is supposed to be a proposition that is true solely in virtue of its constituent concepts and the relations between those concepts. For example, defenders of the view that there are conceptual truths often claim that the proposition that all vixens are foxes is true solely in virtue of the concepts VIXEN, FOX, and UNIVERSAL QUANTIFICATION, and entailments between those concepts. These philosophers also often contrast this example with the proposition that some vixens are hungry, which they claim is not true solely in

### 3. A Working Model of Philosophical Analysis

In the previous sections we got some indication of the range of views that have been advanced on the topic. It is helpful to fix on a particular

<sup>4</sup> See especially Carnap (1950, ch. 1).

virtue of the concepts VIXEN, HUNGRY, and EXISTENTIAL QUANTIFICATION and entailments between those concepts. Now suppose the English word "vixen" expresses the concept *F*, and that the English word "foxes" expresses the concept *G*. Then the English sentence (S)

(S) Vixens are foxes.

expresses the conceptual truth that *F*s are *G*s. Sentence (S) is true in virtue of expressing a conceptual truth. Another way of saying this is that (S) is an **analytic truth** (as opposed to a **synthetic truth**). A conceptual truth is a **conceptually necessary truth**. So (S) expresses a **conceptually necessary truth**. Another way of saying this is that (S) is an **analytically necessary truth**.

(3) *An analysis is informative.* The analysis of the term "*F*" in terms of "*G*" provides novel information about the meaning of "*F*." One way in which this requirement has been met is by further requiring that the analysis is non-circular. This means that if "*F*" is analysed by "*G*," "*G*" is not analysed by "*F*." This non-circularity requirement also establishes the *direction of analysis*: it establishes which term is the *analysandum* and which the *analysans*.

A good illustration of this requirement is the Euthyphro dilemma. Plato's dialogue *Euthyphro* contains a discussion between Socrates and Euthyphro about what it is to be pious. Socrates asks the following question: "Is what is pious loved by the gods because it is pious, or is it pious because it is loved by the gods?" (Plato *Euthyphro* 10a). Socrates' question can be seen as posing a dilemma. Socrates is asking what the relation is between being pious and being loved by the gods. Which one is to be analysed in terms of the other? There are two options: either things are pious because they are loved by the gods, or the gods love pious things because they are pious.

Consider the first option. This says that the analysis of why things are pious is that they are loved by the gods. But unless some reason is given as to why the gods love those things and not other things instead, it seems arbitrary and gratuitous that the gods should love the things that they do. The gods might have loved other things instead. So if we are going to analyse what it is for a thing to be pious by saying that the gods love it, we need to say why the gods love that thing.

Furthermore, what is so special about the gods? Mortals love various things too. Why is a thing's being pious to be analysed in terms of the gods loving it rather than in terms of the mortals loving it? Unless some reason is given as to why there is an analytic connection between

the gods loving a thing and that thing being pious, it seems arbitrary to analyse piety in terms of what the gods love rather than in terms of (say) what mortals love. So we need to show that there is an analytic connection between *the gods loving certain things* and *those things being pious*.

If we were to say that the gods love those things because those things are pious, we would have shifted to the second option. Notice that the two options are mutually exclusive. It cannot be the case both that things are pious because the gods love them *and* that the gods love those things because they are pious. That would be a vicious circle. It says that things are pious because (following the circle around) they are pious. So the analysis of why things are pious cannot be because things are pious.

The second option analyses why the gods love certain things in terms of those things being pious. To avoid circularity, what it is for things to be pious cannot then be analysed in terms of its being loved by the gods. It will have to be analysed in some other way.

The dilemma that Socrates poses then comes to this. The claim that a thing is pious because it is loved by the gods seems to make an arbitrary connection. Why the gods? And why do the gods love those things? On the other hand, if the claim is that the gods love certain things because those things are pious, what it is for a thing to be pious has to be analysed in some independent way — some way that is independent of the claim that the gods love those things.

The Euthyphro dilemma has immediate relevance for the divine command theory of morality. This theory analyses what it is to be morally wrong as what God forbids and what it is to be morally right as what God commands. The dilemma this theory faces is that either things are morally wrong because God forbids them, or God forbids certain things because they are morally wrong. The first limb of the dilemma raises the issue of arbitrariness. Why should God forbid certain things and not others? If God has no reason for forbidding those things, then his ruling seems arbitrary. The second limb of the dilemma analyses why God forbids certain things in terms of those things being morally wrong. To avoid a vicious circle, what it is for something to be morally wrong cannot be analysed in terms of God's forbidding it. To sum up, the dilemma is that either the divine command theory arbitrarily takes what is morally wrong to be what God forbids, or things are not morally wrong because God forbids them.<sup>5</sup>

The general point behind the Euthyphro dilemma is the claim that the relation of analysis is irreflexive: the analysis of a concept *F* cannot involve the same concept. Consequently, it is not possible to analyse *F* in

5 The Euthyphro dilemma remains controversial. See Joyce (2002) for criticism of it.

terms of the concept *G* and analyse *G* in terms of *F*. This is the basis for the charge of vicious circularity made above. Another consequence is that analysis is also an asymmetric relation: if *F* is analysed in terms of *G*, then *G* cannot be analysed in terms of *F*.

Let's now continue with the remaining claims of our working model of philosophical analysis.

(4) *An analysis is knowable a priori*. The biconditional stating the analysis

$$(Vx) (Fx \leftrightarrow Gx)$$

can be known *a priori*. Furthermore, it can be known *a priori* that the biconditional is necessarily true, that it is analytic, and that it provides the analysis of "*F*".

Given certain assumptions about the connection between what is analytic and what is knowable *a priori*, all of these claims can be derived from requirement (2) above. Assuming that an analysis states an analytic truth, and that every analytic truth is knowable *a priori*, it follows that the biconditional stating the analysis can be known *a priori*. An analytic truth is a sentence that is true solely in virtue of its meaning. Given that we can know *a priori* what a sentence means, we can know *a priori* that it is true in virtue of its meaning. If the biconditional is an analytic truth, it is knowable *a priori* that the biconditional is an analytic truth. Lastly, if it is an analytic truth that the biconditional states the analysis of "*F*", then, again given the assumption that every analytic truth is knowable *a priori*, it is knowable *a priori* that the biconditional provides the analysis of "*F*".

(5) *An analysis is testable by the method of hypothetical cases*. This method concerns how philosophical analyses are tested. By considering whether you would be inclined to apply a certain term in various hypothetical situations, you can, without further reflection, discover *a priori* various analytic truths about the necessary or sufficient conditions for the correct application of a term. This is the method of hypothetical cases—of appealing to philosophical intuitions—mentioned in §1.

For example, a theory in the philosophy of mind called "analytical behaviourism" analyses claims about mental states as claims about how people behave and are disposed to behave. On this theory, the claim that Ned believes that it is sunny is to be analysed as the claim that Ned does, or is disposed to do, such things as loosening his collar, putting on sun cream, turning up the air conditioning, and so on. How plausible is that analysis? Suppose that Ned turns out to be a robot controlled by radio by computer scientists. All of his behaviour and his dispositions to behave

are produced by the scientists at their controls. Would you be inclined to say that the thing we have been calling "Ned" believes that it is sunny? Would you be inclined to apply the phrase "believes that it is sunny" to Ned? What you are inclined to say, what your intuitions are, can be enough to give you *a priori* knowledge of what is required for the phrase "believes that it is sunny" to apply to anything.

Intuitions are used to test an analysis in the following way. The analysis says that, necessarily, anything is an *F* if and only if it is a *G*. We imagine a hypothetical case in which something is an *F*. We then see whether we have the intuition to describe that thing as a *G*. If we do, the intuition is some evidence for the analysis being true. If we have the intuition that the thing is not *G*, that intuition is some evidence that the analysis is false. If we have no intuitions, we consider another hypothetical case and proceed as before.

It might be that some intuitions provide evidence for an analysis whereas others provide evidence against the same analysis. Depending on the balance of evidence, various intuitions may need to be revised or discarded. Some intuitions may provide evidence against a given analysis but, because the analysis is especially simple or fruitful or illuminating, the balance of evidence may be against those intuitions. Our earlier hypothetical example of Ned might produce an intuition against analytical behaviourism. But here is another hypothetical example that might produce an intuition *for* it. Suppose that you were to discover that none of your best friends' heads contain brains but instead cogs and motors that regulate their behaviour towards you. Do you have the intuition that this would be a situation in which you discovered that what you took to be your friends never had any mental life at all, or (as analytical behaviourism would claim) that you discovered something about the inner goings on of friends who thought and cared about you all along? Reaching a philosophical analysis will then be a matter of weighing up proposed analyses against intuitions and reaching a reasoned balance between them. This method is known as *reflective equilibrium*.<sup>6</sup>

The above five claims, (1)–(5), characterize a widely held model of philosophical analysis. Claims (1)–(4) describe the nature of philosophical analysis. Claim (5) describes how philosophical analyses are tested. The remainder of this chapter will be concerned with the nature of philosophical analysis, and the next chapter will be concerned with the method of hypothetical cases.

<sup>6</sup> Goodman (1955, 61–66), Rawls (1971, 19–20), DePaul (1998), and Henderson and Horgan (2000, 72–73). But see also Cummins (1998).

Here, in summary form, are the five claims that constitute our working model of philosophical analysis:

- (1) An analysis has the logical form of a universally quantified biconditional:  $(\forall x) (Fx \leftrightarrow Gx)$
- (2) An analysis is necessarily true.
- (3) An analysis is informative.
- (4) An analysis is knowable *a priori*.
- (5) An analysis is testable by the method of hypothetical cases.

It is helpful to have an example or two of how this model is supposed to apply. A simple example in which the model appears to apply is the following analysis of "x is an even number":

$(\forall x)$  (x is an even number  $\leftrightarrow$  x is divisible without remainder by 2)

Another example: Suppose that we want to analyse sentences of the form "x is the same person as y." An analysis in terms of our working model would be:

$(\forall x) (\forall y)$  (x is the same person as y  $\leftrightarrow$  (x is a person and y is a person and  $x = y$ ))

The analysis says that, for any x and y, x is the same person as y if and only if x and y are persons and x is identical to y. This analysis has the form of a universally quantified biconditional. It states necessary and sufficient conditions for a pair of things being the same person. It also is necessarily true: it applies to any possible pair of things. It is at least moderately informative. And we can know it to be true *a priori* by understanding the terms involved, and by considering what we would say about hypothetical cases of things that are the same person or hypothetical cases of things that are people and are identical. This example leaves open how sentences of the form "x is a person" should themselves be analysed. That task faces difficulties of its own.

The tripartite analysis of knowledge also seeks to meet the requirements of our working model. This analysis runs:

$(\forall x) (\forall p)$  (x knows that p  $\leftrightarrow$  x has a justified true belief that p)

Whether this analysis is correct was notably challenged by Gettier (1963).

## 6. The Paradox of Analysis

An important issue that Moore's open question argument raises is a general one about philosophical analysis. A philosophical analysis of a term's

meaning needs to be both true and interesting. If the analysis is uninteresting, then it does not deepen our understanding of the term's meaning. For instance, an analysis that said that "intrinsically good" has the same meaning as "intrinsically good" would be true but uninteresting. But for an analysis to be interesting, it needs to be not obviously true. Ethical naturalism offered an analysis of "intrinsic goodness" that attempts to meet these twin requirements of being true and being interesting. Moore, by contrast, seems to require that an analysis of "intrinsically good" should be both true and obvious. Taken together his requirements rule out the possibility that there is any true and interesting analysis of the meaning of "intrinsically good." But notice that his requirements rule out the possibility that there is any true and interesting analysis of the meaning of any term whatsoever. Suppose, for instance, that we tried to analyse the meaning of "knowledge" by saying that it has the same meaning as "justified true belief." That would be an interesting analysis. But Moore could point out that whereas "knowledge is knowledge" is obvious to anyone who understands it, "knowledge is justified true belief" is not obvious to anyone who understands it. Moore would be committed to rejecting this analysis of the meaning of "knowledge" for just the same reason as he rejected the ethical naturalist's analysis of the meaning of "intrinsically good." More generally, if Moore's open question argument were valid, it would show that the meanings of no terms are analysable (Lewy 1964, 302 and Baldwin 1990, 88). It would follow that philosophical analysis, understood as meaning analysis, is impossible. It then seems that Moore's argument proves too much. A likely diagnosis of where the argument goes wrong is that it erroneously assumes that any analysis, and more generally any conceptual truth, will be recognized as obviously true by anyone who understands it. That assumption seems mistaken. Here is another counterexample. Someone may have the concept GREAT-GRANDPARENT and the concept SECOND COUSIN (i.e., the concept CHILD OF SOMEONE WHO IS A CHILD OF SIBLINGS). Nevertheless, it may not be obvious to them — it may be an open question to them — whether second cousins share two great-grandparents. It may take a little time and reasoning for them to realize that it is a conceptual truth.<sup>13</sup>

Moore's open question argument apparently trades on what has been called the paradox of analysis, a paradox that precisely says that philosophical analysis, understood as meaning analysis, is impossible (Wisdom 1934, 79 and Langford 1942). Langford writes:

Let us call what is to be analyzed the analysandum, and let us call that which does the analyzing the analysans. The analysis then states an appropriate relation of equivalence between the analysandum and the analysans. And the paradox of analysis is to the effect that, if the verbal expression representing the analysandum has the same meaning as the verbal expression representing the analysans, the analysis states a bare identity and is trivial; but if the two verbal expressions do not have the same meaning, the analysis is incorrect. (Langford 1942, 323)

The paradox can be set out a little more fully as follows:

- (1) A philosophical analysis of a concept, *F*, should say what that concept is identical to.
- (2) Suppose that a given philosophical analysis of the concept *F* says that it is identical to the concept *G*.
- (3) So if the concepts *F* and *G* are not identical, the analysis is false. (By (1) and (2)).
- (4) Alternatively, if the concepts *F* and *G* are identical, then what the analysis says (that the concept *F* is identical to the concept *G*) has the same content as the claim that the concept *F* is identical to the concept *F*.
- (5) But it is uninteresting ("trivial," as Langford calls it) that the concept *F* is identical to the concept *F*.
- (6) So if the concepts *F* and *G* are identical, what the analysis says is uninteresting. (By (4) and (5)).
- (7) So either the analysis is false or it is uninteresting. (By (3) and (6)).

Premise (4) can be supported in two ways. Here is the first way. An analysis of the concept *F* says what that concept is identical to, namely the concept *G*. So the analysis says that concept *F* is identical to itself. The claim that concept *F* is identical to concept *F* also says that concept *F* is identical to itself. So that claim says the same as the analysis. Here is the second way. There is a principle of semantics known as the principle of compositionality. This principle says that the meaning of a sentence is determined by the meaning of each of its components (the concepts that they express) and by how they are arranged. Suppose that the analysis is correct and so that concept *F* is identical to concept *G*. Then the claim that concept *F* is identical to concept *G*, and the claim that concept *F* is identical to concept *F*, consist of the same terms arranged in the same

<sup>13</sup> See also Baldwin (1990, 210–11), and Darwell, Gibbard and Railton (1992, 115).

way. By the principle of compositionality those claims have the same meaning — they say the same thing.

The argument (1)–(7) poses a dilemma for any philosophical analysis, where philosophical analysis is understood as a meaning analysis. Either the analysis captures the content of the concept in question, or it does not. If it does not, then the analysis is not informative and so is not interesting. If it does, then the analysis is false. The paradox arises because intuitively it seems possible for philosophical analysis to produce true and interesting analyses of concepts. In fact, the paradox of analysis is doubly paradoxical because the paradox seems to be an interesting result of an analysis of the concept of a philosophical analysis.

A simple attempt at a solution would be metalinguistic (Ackerman 1981, 1990). Consider the following analysis:

(A)  $(\forall x)$  ( $x$  is procrastinates  $\leftrightarrow x$  defers action)

A metalinguistic solution would construe (A) as being about not just the concepts of procrastinating and deferring action, but also about the terms “procrastinates” and “defers action.” (A) would be construed as (B):

(B)  $(\forall x)$  (the concept expressed by “procrastinates” applies to  $x \leftrightarrow$  the concept expressed by “defers action” applies to  $x$ )

A problem with this attempt is that the same analysis can be expressed in different words. For example, assuming that “procrastinate” is synonymous with “postpones things he or she should be doing” (A) can be expressed as:

(C)  $(\forall x)$  ( $x$  is procrastinates  $\leftrightarrow x$  postpones things he or she should be doing)

(C) does not say anything about the term “defers action.” So (C) does not mean the same as (B). But (A) means the same as (C). So (A) does not mean the same as (B) (Rieber 1994, 107–09).

A more promising line of solution would be that in an analysis the term for the analysandum and the term for the analysans should differ syntactically and should have different semantic structures. What is a semantic structure? Let  $T$  be a term and each of  $T_1, \dots, T_n$  be its components. The semantic structure of  $T$  is the property of being a term each of whose components have the respective meanings of  $T_1, \dots, T_n$ . For example, the semantic structure of “female fox” is the property of being a term with

two components whose respective meanings are the property of being female and the property of being a fox.

[Two] expressions have the same semantic structure if and only if they are synonymous and for each meaningful component of one there is a corresponding component of the other which means the same. (Rieber 1994, 110)

It follows that “female fox” and “vixen” have different semantic structures. For each semantic structure there is a corresponding concept. There may be more than one semantic structure corresponding to the same concept. For example, the semantic structures of “female fox” and “vixen” both correspond to the concept FEMALE FOX (i.e., the concept VIXEN). Consequently, semantic structures are more finely individuated than concepts. Different terms can have the same semantic structure. So terms are more finely individuated than semantic structures. Let’s introduce the notation “ $\alpha \approx \beta$ ” to indicate that a term refers to its own semantic structure. Recall (A):

(A)  $(\forall x)$  ( $x$  is procrastinates  $\leftrightarrow x$  defers action)

What (A) says is given by (D):

(D)  $(\forall x)$  (the concept corresponding to «procrastinates» applies to  $x \leftrightarrow$  the concept corresponding to «defers action» applies to  $x$ )

The terms “procrastinates” and “defers action” in (D), and so in (A), refer to their semantic structures. What (A) says can be said using other expressions provided that they refer to the same semantic structures.

How does this apparatus solve the paradox of analysis? A statement of an analysis refers to the semantic structures of the terms flanking the biconditional. Recall premise (4) of the paradox:

(4) If the concepts  $F$  and  $G$  are identical, then what the analysis says (that the concept  $F$  is identical to the concept  $G$ ) has the same content as the claim that the concept  $F$  is identical to the concept  $F$ .

That premise assumes that the analysis is simply about the concept  $F$ . That assumption is rejected by the current suggestion. It takes the analysis to say that concept  $F$ , the concept corresponding to one semantic

structure, is identical to concept *G*, the concept corresponding to another semantic structure.

The argument from compositionality says that the meaning of a sentence is determined by the meanings of its components, and how those components are combined. But compositionality does not apply to expressions that refer to their own semantic structures. “[The] contribution made by an *expression which refers to its own semantic structure* to the meaning of a sentence depends not on its own meaning — or at least not solely on its own meaning — but rather on its semantic structure” (Rieber 1994, 112).

To sum up, the paradox of analysis is a powerful challenge to the model of philosophical analysis that we are working with. Nevertheless, that model can be supplemented with an account of the semantics of sentences stating analyses that offers a promising solution to the paradox.<sup>14</sup>