Interpretation and Understanding

Catherine Z. Elgin

Abstract: To understand a term or other symbol, I argue, it is generally neither necessary nor sufficient to assign it a unique determinate reference. Independent of and prior to investigation, it is frequently indeterminate not only whether a sentence is true, but also what its truth conditions are. Nelson Goodman's discussions of likeness of meaning are deployed to explain how this can be so.

According to a familiar and not unattractive theory of language, to understand a word is to interpret it correctly, the correct interpretation being the one that correlates it with the right referent. The correct interpretation of 'elephant' maps the word onto members of the class of elephants, not the class of bumble bees; the correct interpretation of the name 'Julius Caesar' maps the name onto the emperor, not onto his dog. To understand a denoting symbol is to know what in the world it refers to; and to understand a sentence is to know how things must stand for the sentence to be true. Thus the sentence 'Caesar is an elephant' is true just in case the entity denoted by the term 'Caesar' is in the extension of the predicate 'elephant'. Understanding the sentence does not, of course, require knowing whether it is true. We need only know what it takes for the sentence to be true. The rest is a matter for empirical investigation. But the understanding that the correct interpretation supplies determines precisely what is to be investigated. It specifies exactly what about the world we want to know.

For a sentence like 'Caesar is an elephant', this sort of account (with suitable elaborations to accommodate greater complexity) may seem plausible. We take ourselves to know whom the name 'Caesar' denotes, and what extension the predicate 'elephant' picks out. So we take ourselves to know under what conditions the sentence 'Caesar is an elephant' is true.

The question is whether such knowledge is paradigmatic of linguistic competence in general. I suggest that it is not. I will argue that to understand a word, sentence, or other symbol, it is generally neither necessary nor sufficient to assign it a unique, determinate reference. Our semantic competence consists in knowing both more and less than the familiar account suggests. Reference fixing is often partial and is often the outcome of empirical investigation rather than a prerequisite for it. The result is that the process of settling on an interpretation and arriving at an understanding of the fact(s) it concerns is a constructive interplay, fraught with contingencies. Independent of and prior to investigation, it is frequently indeterminate, not only whether a sentence is true, but also what it would take for the sentence to be true.

To see this, we need not move far from our original example. The sentence

Caesar is an elephant.

seems relatively unproblematic, since we take ourselves to be pretty good at telling what is and what is not in the extension of the term 'elephant'. But consider the sentence

Caesar was a tyrant.

True or false? Well, we're apt to say, it depends on what you mean by 'tyrant'. Under some reasonable interpretations it comes out true; under others, it comes out false; under yet others it may be utterly indeterminate. This seems to be exactly what we should say. But it means that we've been using the word 'tyrant' all these years without a fixed interpretation.

We might hope to evade the embarrassment by retreating to idiolects. Perhaps each speaker assigns a determinate interpretation to the term, but not necessarily the same one. The indeterminacy we've found would then result from

the fact that I assign one interpretation to the word 'tyrant' while you assign another. This is implausible. If each of us consults her own linguistic intuitions, she's apt to discover the same indeterminacy. Whether my sentence is true depends on what *I* mean by 'tyrant', and my intuitions are consonant with my meaning any of several different things.

Moreover, the proposal does not square with the way we acquire language. We learn our native tongue by being brought up in a community where it is used. Through exposure to the language, we learn terms and conditions on application. What we gain is a sense of when members of the community take a term clearly to apply, when they take it clearly not to apply, and when they take it neither clearly to apply nor clearly not to apply. This can be a very nuanced, contextual matter. A child learns fairly early that the word 'tall' is relative. Whether it applies to an object of a given height depends on what sort of thing that object is. What is tall for a house is not tall for a mountain, and what is tall for a man is not tall for a tree. She may need to acquire a whole body of theory to learn even the clear cases of 'molecule' or 'tyrant'. She may learn that different subpopulations take the term to apply in different situations. But because language learning consists accommodating one's verbal behavior to the mores of the linguistic community, a speaker can by this method learn no more than how the terminology is used in that community. When the child uses the term 'tyrant' in the way the rest of the community does, she has mastered the term. If she wants to know anything more about tyrants, she will have to study politics, not language. Whatever is indeterminate in the verbal behavior of the linguistic community is not something she can learn by modelling her behavior on the behavior of members of that community. If differences in interpretation are not reflected in linguistic practice,

3

the language learner has no access to them. If they are, they afford her reason to believe that the term is vague, ambiguous, or contextual. In none of these cases, does she have a basis for assigning a unique interpretation, possibly a different one from that assigned by her interlocutors.

Language learning involves more than just learning what various terms apply to. We also learn the community's criteria of application -- we learn, that is, how they tell whether a given term applies. Things that are appropriately like the clear cases of the term are also instances of it. The question, of course, is what makes for appropriate likeness. Learning that involves an understanding of what sort of term we are dealing with, what sort of role it plays in the language. If the word 'tiger' is a species term, then things appropriately like tigers are members of the same species. Deciding what kind of a term it is is determining how to project beyond the clear cases.

We glean this information too by becoming members of the community where the terminology is used. But the community need not speak with one voice. Hilary Putnam [1975] and Tyler Burge [1979] have argued convincingly that there is a division of linguistic labor. My competence with terms like 'arthritis' and 'elm tree' derives from my place in a linguistic community, some of whose members know what they are talking about. The rest of us rely on and defer to the experts. Arguably the same holds for 'tyrant'.

What constitutes expertise? Physicians who are experts on arthritis know a lot about the ailment: its causes, manifestations, predisposing conditions, treatments, complications, and the like. They also know a lot about how to study the ailment. Experts on elm trees and tyrants have similar knowledge about their respective fields. The details don't matter. The point is that experts tend to have a

4

relatively wide and deep understanding of their subject matters and how to investigate them. But this understanding does not require, and is not limited to, knowing the extension of the term that grounds the expertise. Experts on arthritis, although they are vastly more knowledgeable than laymen, do not know and do not purport to know exactly what arthritis is, nor do they know or purport to know exactly which members of the population have the ailment. Experts on tyrants do not know and do not purport to know exactly what constitutes a tyrant, hence do not know and do not purport to know exactly who belongs to the extension of the term.

The experts, moreover, do not always agree. Perhaps medical opinion is divided about whether arthritis is a syndrome or a disease. In that case, opinion is divided about the basis on which to project from clear to unclear cases. Even among those who consider it a syndrome, there may be differences of opinion about exactly what symptoms are part of the syndrome. Again, opinion is divided about how to project from clear to unclear cases, hence about what disputed cases are in fact cases of arthritis.

Still, one might argue, it doesn't follow that the terms in question don't have a unique, determinate extension. Opinion is now divided about what is and what is not a case of arthritis, or about who is and who is not a tyrant. But at the end of inquiry, one might urge, the question will be answered. To assume that the interpretations of our current terms are now fixed, we need only recognize that the requisite linguistic community is extended in time. So our descendants' discoveries figure in the determination of what we are now talking about. In effect, the experts we defer to are the intellectual heirs of the current experts.

But we should not assume that the future course of inquiry is somehow laid

out in advance. There is typically no unique trajectory from what we now take ourselves to know about a subject to the verdict that will be reached at the end of inquiry. What our intellectual heirs will conclude is affected by choices we make. Rather than an antecedently fixed extension determining what current inquiry is looking for, what current inquiry is looking for and finding affects the further specification of the extension of our terms. Where inquiry will end up depends on the choices inquirers make. And a variety of choices may be equally reasonable.

Opinion is currently divided, perhaps, about whether arthritis is a syndrome or a disease. But scientists are working on the matter, and eventually will find the answer. What will they find? One possibility is this: They find that lots of the cases we currently consider clear cases of arthritis have the same cause. This discovery affords some incentive to opt for the disease model and say that all and only conditions with that cause are genuine cases of arthritis. This requires rejecting as instances of arthritis some of the cases that were previously considered clear. Even if they exhibit the same symptoms and respond to the same treatment, once we have opted for the disease model, they are to be excluded, for they lack the now requisite etiology. It also requires saying that any condition with the same etiology whether or not it exhibits the same (or indeed any) symptoms, constitutes a case of arthritis. If we want a causal story, these are prices worth paying. But notice the 'if' clause. What justifies the decision to exclude some conditions from and include other conditions in the extension of the term 'arthritis' is the desire to have the term play a particular conceptual role. This is entirely reasonable. But it is not mandatory. Even if lots of the cases we currently consider clear cases of arthritis have the same cause, there might be reason to prefer a characterization that treats arthritis as a syndrome -- a constellation of symptoms. If, for example, our interest

is ergonomic, it may be relatively unimportant why people come to have these symptoms. The critical fact is that they do. We want to delineate the class of individuals who have a particular cluster of symptoms so that we can construct devices and develop techniques to minimize their discomfort and disability. If we or our descendants decide to treat 'arthritis' as a syndrome concept, all and only conditions with the same symptoms will be in the extension of the term. Again, some currently accepted cases may be thrown out and currently disputed or rejected cases included in the extension of the term. But it will not be the same cases as under the etiological account.

The question of how to revise and refine current usage depends on what we want our concepts to do. Currently accepted usage constrains, but does not determine the answer to this question. The choices we make about how to revise and refine, and how to decide among reasonable aims, affect the constitution of the individuals and kinds our descendants will recognize. The course of inquiry is not the straight and narrow path to the ultimate Truth. Rather, decisions we make, options we choose, affect the future direction of inquiry. And there are some points where a variety of options, pointing in different directions, are equally good on balance. What truths we will find depend on what direction we choose to go.

So far, the constraints we've recognized have all been factual. They derive from the ways terms have been used and deemed acceptable in statements of fact, and the interests those statements of fact are designed to serve. But we do other things with our terminology besides state what we take to be facts. We contrive fictions as well. And the way we use our language in fiction feeds back on what we are willing to say about matters of fact.

According to Nelson Goodman [1949, 1953], at term's meaning is a function

of its primary and secondary extensions -- of the things the term denotes and the things compounds containing the term denote. Its applications to matters of fact constitute its primary extension. Its occurrences in fiction contribute to its secondary extension. Consider the word 'dog'. Its primary extension consists of all the dogs. Its secondary extension contains dog-stories, dog-pictures, and the like. Some dogs answer to no dog-description; some dog-descriptions describe no dog. But even fictive dog-descriptions and dog-pictures belong to the secondary extension of, hence figure in the meaning of, the word 'dog'. If Goodman is right, the stories we tell, and the pictures we paint, affect the meanings of the words we use. If I'm right, they also affect the choices we make about which extensions to assign to our terms.

We saw that language learning involves learning to extrapolate from clear to unclear cases. Considerations of purpose and context influence our extrapolations, but they rarely suffice to determine just what lines to draw or precisely where to draw them. We know that poodles and spaniels and the like fall under the predicate 'dog'. Precedent suffices to enable us to extrapolate painlessly to setters and Dalmatians. But it does not tell us, for example, whether to include coyotes in the extension of the term 'dog'. Coyotes are in some respects like and in some respects different from the animals we have no qualms about considering dogs. The question is whether the similarities or the differences loom larger. Studying the clear instances and clear counter-instances of a predicate does not tell us where to draw the line.

Both primary and secondary applications afford avenues of extrapolation. From accepted instances, we extrapolate to further instances, thus augmenting primary extension. Once we've accepted the extrapolation, we have a new

precedent class that includes poodles, spaniels, setters and Dalmatians. Henceforth, all count as clear cases. Further candidates for admission should be like the members of the augmented class. Moreover, from accepted instance-portrayals -- descriptions, pictures, and the like -- we extrapolate to augment secondary extension. Thus on the basis of the pictures and descriptions we have already classified as dog-portrayals, we generalize to further pictures and descriptions that we are willing to count as such. In both cases, each new application alters the precedent class against which further candidates will be judged. Even if our original paradigm provided little incentive to classify Newfoundlands as dogs, once we have so classified St. Bernards, leaving Newfoundlands out would seem arbitrary. And even if our paradigm descriptions provided little incentive to classify the cartoon figure Snoopy as a dog-picture, once we have so characterized Goofy, we have no grounds for excluding Snoopy. So things we were initially indifferent about may come to be naturally included or naturally excluded as precedent evolves.

Goodman emphasizes the logical independence of primary and secondary extensions. But, I believe, much of their cognitive significance derives from their interanimation [Elgin, 1997]. New factual applications provide fodder for fiction. And new fictional applications influence further findings of fact. It is hardly surprising that what we count as a dog influences what we count as a dog-description. But the converse may be unexpected. What we proffer and accept as dog-descriptions and dog-pictures influences what we count as dogs.

The primary extension of the term 'dog' affords no rationale for drawing the boundaries where we do. Coyotes, for example, differ little from some of the canines we keep as pets. But fictional dog-stories portray their subjects as thoroughly domesticated -- as loyal, loving companions and/or devoted servants.

9

Even stories about vicious dogs construe their subjects as exceptions to the norm. Familiar fictions thus disincline us to extend the term 'dog' to coyotes, jackals, and other notably nasty members of the canine family. No one raised on stories like 'Lassie' or 'Rin Tin Tin' would feel comfortable calling such animals dogs. Animals that do not differ much from dogs are excluded from the term's extension because they do not conform to the dog-descriptions favored by our fictions. The residues of fiction thus infuse findings of fact. For literal extrapolations from the cases we consider clear are sensitive to secondary extensions. Meaning is not a fixed feature of terminology, given once and for all with the mastery of our words. It is a dynamic interplay, evolving over time in response to the facts we discover and the fictions we contrive.

It is not only fictional secondary extensions that have this feedback effect. Factual ones do too. Among the secondary extensions of 'arthritis' are arthritis-descriptions that occur in medical journals, medical records, complaints of patients, and so on. These are all statements of fact that are, and should be taken seriously by the medical community. Many of them presumably are true. But not all facts get stated. There are likely to be so-called silent cases of the ailment -- cases that never get labeled 'arthritis'. And there are aspects of the medical conditions of patients who are diagnosed as arthritic that are deemed unworthy of mention. This is unsurprising. Any description is selective. The point is that the selection has consequences for future findings of fact. The features that appear in the arthritis-descriptions generated and accepted by the medical community become salient. Hence they are apt to become distinguishing marks of the condition. Hitherto borderline cases get settled through having or lacking the features that the accepted descriptions have highlighted. The critical question then is not the way

10

the world is *simpliciter,* but, as Goodman says, the way the world has been described and anticipated in words [Goodman, 1960].

I said earlier that the correct answer to the question, 'Was Caesar a tyrant?' is 'It depends'. Now it seems that the correct answer to the question, 'Did Caesar suffer from arthritis?' might also be 'It depends'. In both cases, it depends not just on historical information that we are not privy to, but on the future course of inquiry, which is affected by choices we and our descendants make. This has the uncomfortable sound of a claim that the future can change the past.

Caesar is dead. His political demeanor and physical constitution were whatever they were, we want to say. Possibly during his lifetime they could have been influenced by the actions and choices (and maybe even the descriptions) of others. But there's no doing anything about them at this late date. To contend otherwise is crazy.

Well, yes and no. The contention that the extensions of the terms 'arthritis' and 'tyrant' are indeterminate is the contention that the commitments we have about their instances are insufficient to fix a unique extension for each of the terms. Any of a number of intersecting extensions satisfies all the conditions we are prepared in good conscience to put on the application of the terms. For any one of the candidate extensions, it is determinate whether Caesar belongs. But because it is indeterminate which candidate is *the* extension of the word 'tyrant', if Caesar is among the disputed cases, it is indeterminate whether he was a tyrant. And because as inquiry proceeds and new descriptions are contrived and accepted, we refine our categories and narrow the range of disputed cases, it would not be surprising if our descendants are in a position, as we are not, to yield a definite verdict. This would not require them to settle on a unique extension for the term,

only to narrow the range of options enough that Caesar is definitely in or definitely out. Future refinements of categories can and often do affect what we are in a position to say about the past.

   If I'm right, interpretation is not a matter of determining the uniquely correct referent of a symbol. It is holistic-- depending on how the symbol functions in context, what is presupposed about its primary and secondary extensions and its linguistic and extra-linguistic milieu. Quine has long denied that there is clear distinction between matters of language and matters of fact [see Quine, 1953]. Nor, as we've seen, is there a clear distinction between interpretation of language and understanding of facts.   Interpretation and understanding inextricably intertwine. And their deliverances are apt to be open-ended.   Previously accepted usage supplies precedents. But available precedents do not always determine how to go on. Where they do not, we face a choice. Although the choice is influenced by interests and objectives, these are not always sufficient to settle the matter. Where they are not, the choice is arbitrary. But, arbitrary or not, our choices about such issues affect the constitution of the precedent class against which future cases will be tested. They affect what, henceforth, is to count as a tyrant, a case of arthritis, a dog. By the choices we make, then we construct the categories that fix the fact that Caesar was, or that he was not, a tyrant; the fact that arthritis is, or that it is not, a disease; the fact that a coyote is, or that it is not, a dog. In so doing, we participate in the construction of the world that we and our descendants will inhabit.*

<center>Notes</center>

*This paper was presented at the conference *Weisen der Welterzeugung,* at the

Institüt für systematische Forschung, in Heidelberg in 1998, and will be published in German in the proceedings of that Cogress entitled *Weisen der Welterzeugung,* ed. Hans Rudi Fischer, Carl Auer Systeme Verlag.

References

Burge, Tyler: 1979, 'Individualism and the Mental,' *Midwest Studies in Philosophy,* 4, 73-121.

Elgin, Catherine Z.: 1997: *Considered Judgment,* Princeton University Press, Princeton, pp. 178-195.

Goodman, Nelson: 1949, 'On Likeness of Meaning,' *Analysis*, 10, pp. 1-7.

Goodman, Nelson: 1953, 'On Some Differences About Meaning,' *Analysis* 13, pp. 90-96.

Goodman, Nelson: 1960 'The Way the World Is,' *Review of Metaphysics,* 14, pp. 48-56.

Putnam, Hilary: 1975, 'The Meaning of "Meaning",' *Mind, Language, and Reality,* Cambridge University Press, Cambridge, pp. 227-229.

Quine, W. V., 1953: `Two Dogmas of Empiricism,' *From a Logical Point of View,* Harvard University Press, Cambridge, pp. 20-47.

Catherine Z. Elgin

Graduate School of Education

Harvard University

Cambridge, Massachusetts 02138 USA