

DOES 'RABBIT' REFER TO RABBITS?

MICHAEL LISTON

Associate Professor of Philosophy
Department of Philosophy
University of Wisconsin/Milwaukee

ABSTRACT

It is commonly presupposed that all instances of the deflationary reference schema '*F*' applies to *x* if and only if *x* is '*F*' are correct. This paper argues, mainly on the basis of concrete example, that we have little reason to be confident about this presupposition: our tendency to believe the instances is based on local successes that may not be globally extendible. There is a problem of semantic projection, I argue, and standard accounts that would resolve or dissolve the problem are problematic.

Key words: the problem of semantic projection, reference schema, the problem of induction

1.

We are finite beings who inhabit an idiosyncratic region of the universe. Yet we naturally tend to have beliefs whose content greatly outstrips their parochial, testable consequences. The tension between these observations – the one encouraging modesty, the other suggesting hubris – is well articulated in Hume's discussion of induction. How can we devise an epistemic policy that will reliably carry us from our limited, parochial, and possibly unrepresentative evidence to more universal beliefs? Hume answered that we cannot; any policy we devise will beg the question and depend on a principle of uniformity of nature, a principle guided more by hubris than by modesty. There is nearly unanimous consensus that we cannot answer Hume's problem. We have become modest and learned to live with it without solving it. There is a related problem whose normal form is very close to that of Hume's problem and where we are more inclined to be driven by hubris. The human sensory and cognitive apparatus is a fortu-

itous product of evolution, and we occupy a relatively unusual location in the universe. This should encourage us to be modest about our abilities initially to classify objects and fix our references in an unproblematic way. We should expect our initial, pre-theoretical classifications to be sufficiently parochial and domain-specific that they will not be readily extendible in an unproblematic way. Moreover, human language is itself a fortuitous product of evolution with lots of borrowed and shifted meanings, both noticed and unnoticed. So we should also expect pre-theoretical classifications and reference fixings to bend and wind in ways that are likely to run us into problems when we encounter new objects or attempt new tasks. Finally, our scientific classifications evolve out of our pre-theoretical classifications, and even though their evolution is more guided by critical, self-conscious reflection, it is unreasonable to expect them to be guided by an invisible hand that will remove all the contingency they inherit from their pre-theoretical ancestors and pick up on their own. The foregoing suggests that we should be modest about what our classificatory predicates mean. But we are not. We naturally tend to formulate theories and hypotheses whose content greatly outstrips the parochial domain where our concepts might be expected to have unproblematic application and to consider them as unproblematic. Some of these theories – in philosophy, psychology, and linguistics – concern the very phenomenon in which we are interested: they provide accounts of content. But, given the modest origins of our thought and talk, an account of content should not simply assume that those objects that apparently have unproblematic content – beliefs, sentences, concepts, predicates – when expressed, asserted, or applied in our limited, parochial environment will also have unproblematic content when applied more generally. I will refer to this problem as “the problem of semantic projection” (PSP).

I begin by saying a little more about what I intend the problem to be. We know we are disposed to make mistakes, to mistake horses for cows on a dark night, e.g. But such mistakes can not be comprehensibly expressed, unless ‘horse’ applies to an object if and only if it is a horse. Part of our typical understanding of the claim that we are mistaken is that ‘horse’ applies to horses and not cows. We are extremely confident of all instances of the (referential) schema:

(RS) ‘*F*’ applies to *x* if and only if *x* is *F*

In its most elementary form, (PSP) amounts to the question: is our confidence in (RS) justified? Given the modest origins of our thought and talk, are we justified in our confidence that, when a predicate is applied locally without problems, it will also have unproblematic application conditions across the board?

Some will be immediately inclined to respond that the alleged problem is just the product of confused thinking, for how could it not be the case that ‘horse’ applies only to horses? The mistake mentioned above is not a semantic mistake; it is simply that horses are not cows. Failure of an instance of (RS) would require some predicate’s extension both to include and to exclude some object. And this is not a possibility we can coher-

ently express, since its expression would require 'Fx & ~Fx' to be true of some object. While I will have more to say about this line of thought in section 3, it does initially push us to make the problem more concrete. So, let us look at an example of the problem.

The first American English-speaking colonists arrived in the New World with a stock of classificatory predicates that had unproblematic applications in the Old World. Their common linguistic training disposed them to apply many of these predicates to objects in their new unfamiliar environment. Thus 'rabbit', 'blackbird', 'oriole', 'robin' – and the list is quite lengthy – all came to be used to apply to objects that in some sense are not in the extension of the predicate as it was used in the Old World.¹ Old World usage of 'rabbit' applied to members of the genus *Oryctolagus*; New World usage applied it to members of the genus *Lepus*. New World rabbits are hares, not rabbits. Similarly, New World blackbirds are Icterids, not thrushes; New World orioles are Icterids, not passerines; New World robins are members of *Turdus Migratorius*, not of *Erithacus rubecula*. Suppose we now take an instance of (RS):

(1) 'Rabbit' applies to x if and only if x is a rabbit

The English-speaking colonists (circa 1620) are as disposed to hold this true as our objector of a couple of paragraphs back: its falsity could not be coherently expressed in their language, they might say, for how could 'rabbit' be properly applied to an object that is not a rabbit? They are also disposed to agree that 'rabbit' correctly applies to the various leporidae (all of them hares) they encounter in their new environs. But we are inclined to think the colonists confidently made a natural and understandable error: based on their local Old World successes with classificatory predicates, they over-generalized the application conditions for some of those classificatory predicates. But if the colonists were prone to this kind of error, what makes us think we are not equally prone to it? After all, if we drop the restriction to 1620 English, (1) is just an instance of (RS), instances of which we are inclined to accept so confidently that we think it incoherent even to question them. Similar confidence on the part of the colonists would have been misplaced. So, why should *we* have such confidence in the instances of (RS)? Just as the semantic confidence of the colonists needs to be tempered by the thought that unbeknownst to them their linguistic classifications shifted, so too we should be modest. This is roughly what I intend by (PSP).²

¹ See Mencken 1963, Ch. III.

² My discussion of (PSP) is closely related to, indebted to, and inspired by Wilson 1982 and Wilson 2006. The 'rabbit' example is modelled on his Druid example in Wilson 1982. The themes investigated here also have connections with Kripke's discussions of rule-following in Kripke 1982 and Goodman's discussion of projectibility in Goodman 1965. Whereas Kripke argues that there are no facts about speakers that determine what they mean, I argue more weakly, using different considerations, that we have no reason to be confident that such facts exist. Goodman's answer to the question, why do we project 'All emeralds are green' over 'All emeralds are grue?', in terms of 'green' being an entrenched predicate seems to presuppose an answer to my question, 'how do we know that 'green' applies to green things and not to grue things.'

At this point, many readers will be inclined to respond in one of two ways: either we can sketch a semantic theory that handles PSP adequately, or PSP results from confused thinking. In what follows, I hope to strengthen the argument for the claim that there is a problem, to argue that standard versions of the resolution (section 2) and dissolution strategies (section 3) do not adequately address it, and to suggest some general morals (section 4). The rabbit example is quite simple, but *prima facie* it poses challenges for both resolutions and dissolutions of (PSP).

2.

Many theories in philosophy of language proceed from the assumption that there is some class of facts that underpin our confidence in the instances of (RS). Generally, they postulate an intermediary meaning, that (a) linguistically competent speakers of a language possess with respect to a predicate and (b) determines the predicate's global extension. The idea is that there are meaning-facts about competent speakers that determine the extensions of predicates in their language. If this kind of account could be made to work, we would have an answer to (PSP). Our confidence in our semantic projections would be underwritten by facts about the determination of extension. The problem is to elaborate a plausible account of such meaning-facts that holds any promise of securing the desired outcome.

Classical theories of meaning provide the most elegant account along these lines. Meaningful sentences are well-formed compositions of meaningful syntactic constituents; a constituent is meaningful if it is associated in the mind of the speaker with a (logical complex of) simple concept(s). Simple concepts in turn primitively determine their extensions (at all possible worlds). Thus, when we express any hypothesis, it is meaningful provided it is backed by concepts that we grasp and that determine the truth condition of the hypothesis (at all possible worlds). In particular,

(2) 'Rabbit' applies to *x* if and only if *x* falls under RABBIT

where RABBIT is the concept a linguistically competent speaker associates with 'rabbit'. The speaker grasps the concept and associates it with the predicate; the concept primitively applies to all possible objects of thought; by association the predicate derivatively picks up its extension. Instances of (RS) are not only true, but necessarily true.

But what makes a concept suitable for the role it must play? Various candidates are well known. Empiricists from Hume to Russell often made suitability a function of the "look" or "feel" of an idea as it is presented to our mind. Reacting to the psychologism inherent in the ideas approach, Frege made his concepts abstract and left our grasp of them more or less unexplained.³ For other classical theorists, primitive concepts and

³ Frege was cognizant of the vagaries of natural language. In scientific contexts (where truth and clarity are important), he advocated replacing natural languages by logically perfect languages, where each predicate symbol would

our grasp of them are best thought of as indices of humans' ability to apply or refuse to apply a predicate in response to a presented object.⁴

The problems are well known and do not require rehearsal. For our purposes, the main problem is that they presuppose that the meaning-facts will secure uniform extensions across the board. Most classical theorists thought of their meanings as primitively fixing extensions: they just determine their extensions and truth conditions on their own. But if we are inclined toward modesty and to question whether meanings that are locally unproblematic will work across the board, the classical theories do little to alleviate our worry. We are given no reason to think that empiricist ideas, Fregean concepts, or speakers' capacities to apply predicates correctly will globally extend from the localities where they might be presumed to be in place.

In order to make the foregoing remarks less schematic, let us return to our 1620s English-speaking colonists. There seems to be little warrant to attribute to them a concept – whether we understand concepts as ideas, Fregean concepts, or capacities – that determines a global extension for 'rabbit'. Consider as candidates:

- (2) 'Rabbit' applies to x if and only if x falls under RABBIT
- (3) 'Rabbit' applies to x if and only if x falls under HARE
- (4) 'Rabbit' applies to x if and only if x falls under HARBIT

where RABBIT, HARE, and HARBIT determine respectively $\{x \mid x \text{ is a rabbit}\}$, $\{x \mid x \text{ is a hare}\}$, and $\{x \mid x \text{ is in the Old World and } x \text{ is a rabbit, or } x \text{ is not in the Old World and } x \text{ is a hare}\}$. These assignments are pairwise mutually exclusive. Old World rabbits or New World hares prevent joint use of (2) and (3); New World hares prevent joint use of (2) and (4); Old World rabbits prevent joint use of (3) and (4). Clearly at most one assignment can be correct. But severally, each runs into problems. (2) makes most of their beliefs about New World rabbits false. In Virginia they will say things like, 'There's a rabbit, and he and his family live in a burrow and his young are born blind and hairless'. Under (2) they will have expressed a false belief, since the local leporidae are hares who do not live in burrows and whose young are born fully furred and with eyes open. (3) makes most of their beliefs about Old World rabbits false; e.g., 'The rabbits back home also live in burrows and their young are born blind and hairless'. Attribution of too many false beliefs to them is likely to disconnect their language from the world so badly that it will be hard to make sense of the practical activities (teaching their children to hunt, e.g.) they use language to engage in. Nor can we make sense of their linguistic activities if we postulate a conceptual change leading to a reference change by saying: the backing concept for 'rabbit' shifted from RABBIT before the voyage to

have an extension-determining backing concept. But it is not clear why such languages, when applied to natural objects, will not also succumb to (PSP).

⁴ In some of Schlick's and Carnap's thinking, concepts seem best thought of in this manner.

HARE after the voyage. In their post-voyage usage, they will employ predicates that apply to Old World objects and to New World objects; e.g., ‘The rabbits around here are less edible than the rabbits back in England’. Under the hypothesis of meaning change we are considering, what they are saying is that New World hares are less edible than Old World hares. But that is not what they are saying; what they are saying is that New World hares are less edible than Old World rabbits. All such sentences will get assigned incorrect truth conditions under this meaning change hypothesis.

By contrast, (4) seems to overcompensate for the deficiencies of (2) and (3). It anachronistically and implausibly attributes to them a concept that will induce truth conditions making all the “right” beliefs come out true. Since the colonists systematically apply ‘rabbit’ to rabbits in the Old World and to hares in the New World, assigning HARBIT as the backing concept for ‘rabbit’ will produce truth conditions that make true the sentences they hold true. Two proposals are tempting. One is to say that the backing concept for ‘rabbit’ changed from RABBIT to HARBIT. The problem is that, when they arrive in the New World, their linguistic competence is not especially different from their stay-at-home fellows. They, their Old World past selves, and their Old World current cousins associate the same ideas and concepts with, and possess the same capacities with respect to, leporidae. The other is to say that HARBIT is the backing concept for both Old World and New World uses of ‘rabbit’. Before their voyage, they meant HARBIT by ‘rabbit’ all along. But, this attribution seems odd given that many of them had ideas, concepts, and capacities that enabled them to distinguish Old World rabbits from Old World hares – both are to be found, and were linguistically distinguished as ‘rabbit’ and ‘hare’, in the Old World. This would force us to say that the backing concept for ‘hare’ changed also. Moreover, claiming they meant HARBIT by ‘rabbit’ amounts to claiming rather mysteriously that they were linguistically prepared for an eventuality that was completely outside their ken. Without an underlying assumption of determinacy, the claim seems implausible.

Stripped of such assumptions, we can say what happened. The colonists found it natural to apply ‘rabbit’ to the New World hares they encountered upon arriving, because rabbits are abundant (as pest, food, and pet) in the Old World but non-existent (circa 1620) in the New World and superficially hares look similar to rabbits. In applying ‘rabbit’ to New World hares, the colonists simply exploited superficial similarities and an available linguistic slot, ‘rabbit’, and they need not have been aware of any shift in their linguistic practices. Had the history of their encounters with New World hares gone a bit differently, they might just as naturally have applied ‘hare’ to New World hares. In the Old World, there is a hunting practice where the hunter sends trained ferrets into rabbit warrens to flush out the rabbits. Suppose that, upon setting up in the New World, they immediately went hunting for leporidae with trained ferrets they had brought on the voyage. Realizing that the local leporidae were not burrowers, they would not have called them ‘rabbits’. Instead, they might have exploited another available linguistic slot, ‘hare’, to classify the local leporidae.

We should conclude that no extension-determining meaning-fact of a classical variety will support a claim of uniformity of meaning with respect to the colonists' 'rabbit' predications. No fact involving their ideas, concepts, or capacities is a plausible candidate to determine a global extension for their word 'rabbit'. Because of their similar brains and linguistic and extralinguistic training, the colonists are disposed to agree on their classifications: in the Old World they call rabbits 'rabbits' and hares 'hares'; and in the New World, they call them 'rabbits' or 'hares', depending on how they first encounter them. But this agreement in linguistic practices should not encourage us to think that some underlying fact about conceptual meaning determines the extension of their word 'rabbit', in the sense that a Laplacean genius could predict how they will come to classify New World leporidae purely on the basis of full knowledge of their concepts and their Old World linguistic classifications.

Most philosophers of language nowadays reject the determination of extension by classical concepts. Externalist theorists (Kripke, Putnam, Burge) argue that classical accounts fail, because all assume, in one form or other, two theses that are not cotenable. First, meaning is something competent speakers grasp *in isolation*, whether we gloss this as explicitly having in one's head a classical concept or idea or as having the discriminating capacity to see that a certain condition obtains. Second, meaning determines extension. Putnam's Twin Earth thought experiments are supposed to show that nothing can play both of these roles. Two speakers could share the same concepts, ideas, or discriminating capacities with respect to a predicate, 'is gold', yet intuitively the predicate could possess different extensions (gold and gold*) in the usages of the speakers, given their respective linguistic communities and surrounding environments.⁵ Though they share the classical requirements for applying the term correctly, the term has different extensions in their usage.

Instead, externalist meanings determine extensions. Externalists – here I follow Putnam (Putnam 1975) – hold on to the idea that there are meaning-facts about competent speakers that determine the extensions of predicates in their language. But a competent speaker's knowledge of meaning no longer consists in her grasp of a concept or her possession of a capacity that can be understood in a framework in which she is isolated. It is instead a capacity that presupposes complicated interactions with other language users and with the world. According to externalist accounts, what determines the extension of my term 'gold' is my being a linguistically competent member of a linguistic community that (a) conducts its transactions subject to a division of linguistic labor (DLL) and (b) has its 'gold' usage tied to (the natural kind) gold. I am unable to distinguish reliably gold from non-gold. Nevertheless, I know the meaning of 'gold' and can use 'gold' correctly to apply to gold, because I belong to a community that (a) includes experts who can distinguish gold from non-gold and to whom I'm disposed to defer and (b) has its word 'gold' causally connected to gold. Thus, when we express any

⁵ '*' attached to a word w means 'Twin-w'.

hypothesis, it is meaningful provided it is expressed in the right kind of linguistic and extralinguistic setting, a setting that determines the truth conditions of the hypothesis (at all possible worlds).

The DLL condition (a) is insufficient for determination of extension. If anything underwrites the putative fact that I and my community mean gold by 'gold', it is the word-object regularity (b). To see that DLL is not sufficient, we need only slightly modify Putnam's Twin Earth example (Donnellan 1983, pp. 98-101). Suppose that on Earth and Earth* linguistic practices are identical. All vernacular uses of 'gold' and all expert practices are the same. On both planets there are chemists who use atomic numbering classifications and nuclear scientists who use atomic mass numbering classifications. However, samples of the stuff to which 'gold' is applied on Earth are mainly composed of the only natural isotope gold has here, Au 79/197, whereas samples of the stuff to which 'gold' is applied on Earth* are mainly composed of Au 79/195, the only isotope which naturally occurs there. This seems to show that DLL is no more able to determine reference than are classical concepts. There are no differences in vernacular or expert practices, yet the reference of 'gold' on Earth is Au 79/197, which is not the same as Au 79/195, the reference of 'gold' on Earth*.

If DLL does not suffice to determine the extensions of our predicates, the correlation between our words and objects in the environment does all the work in the Twin Earth arguments in persuading us that extensions are not fixed by classical concepts. But those word-world regularities are assumed rather than worked for. To see what determines the extension of 'gold', we consult our intuitions about what Earth and Earth* speakers are referring to. We learn only that the extension is not determined by any classical concept or, if I am right, by the community's linguistic practices either. But the question of what positively does determine the extensions is left unanswered. The point here is not that there is anything wrong with assumptions like: Earthlings use 'gold' to refer to gold, Earthlings* do not; in the New World, 'rabbit' is used to refer to hares, in the Old World, it is not. But these are just local claims – restricted to certain locales. We have been given no account that would justify extending them across the board. Local semantic ties between words and objects are again simply being assumed to be globally extendible, and this seems to presuppose that some fact of uniformity of meaning underlies our semantic projections rather than to argue for such an underwriting. Again, we have made little progress responding to (PSP).

In Putnam's case, it might be thought that the doctrine of natural kinds can underwrite the needed global extension. Suppose we use 'gold' as a natural kind term and there is a connection between our word 'gold' and objects that are locally accessible to us – gold on Earth. If 'gold' is used as a natural kind term, then it applies to all and only objects of a given kind. Take any object with which we have not interacted. Either it is an object of the same kind as those objects that we have interacted with and call 'gold' or it is not. If it is, it falls in the global extension of our word 'gold'; if it is not, it falls in the

global counter-extension of the word. In this way, local word-world connections can determine global extensions. All that needs to be in place for semantic projection to succeed is a local practice of applying the word to objects that are *in fact* objects of the same kind. Nature takes care of the global determination of extension.

Many philosophers, including a later Putnam himself, are sceptical of the existence of natural kinds. (See *infra* section 3.) Let us suppose for now that there are natural kinds. Let us also suppose that there is in place a practice of applying a word to objects of a certain kind. There is still a problem. First, the notion of a practice of applying a word to objects of a given kind is ambiguous. It may mean that the word is in fact applied in the practice only to objects that are of the same unique kind. Or it may mean that the word is used as a natural kind term in the sense that the speakers intend to use it to apply to all and only objects of a certain kind and an adequate account of their inferential practices would semantically categorize the word accordingly. The first reading simply assumes that some uniformity of meaning obtains that secures global extensions, and thus will not do as an answer to (PSP). The second reading does not assume this. A word may be used as a natural kind term but unbeknownst to its users it may not apply to objects that are of *the same unique* natural kind or of *any* natural kind.

We can easily enough imagine two different evolutions of the language of chemistry. On one, elementhood is tied to chemical combination properties, which depend on atomic number (the number of protons). On the other, elementhood is tied to physical mass properties, which depend on atomic weight (the number of nucleons). Although the language of chemistry took the former path, this was hardly inevitable. There are advantages to proceeding either way: atomic number classification is important for the study of chemical properties; atomic weight classification is important for the study of radioactive decay. It seems likely that the former path was taken because of purely contingent, historical factors. People found that alchemical lore had some practical consequences they could use to their advantage in improving dyeing, brewing, assaying, and other techniques. Theorists (many of whom were also interested in practical applications) got interested. The phlogistonians set the stage for Lavoisier's comprehensive study of chemical combination properties. From Lavoisier's time onward the study of chemical combination properties was intensely pursued, and systematic atomic tables were constructed in the 19th century, most successfully by Mendeleev in 1897. Had alchemy or the practical benefits of chemistry not evolved as they did, or had Rutherford's atom busting experiments occurred 100 years earlier so that scientists became more interested in atomic masses and radioactive phenomena than they were in chemistry, one could well imagine the latter evolution having occurred. And we would now use 'gold' to refer to the element with atomic weight 197. Locally, the two natural extensions are the same: Au 197 is the only natural stable isotope of gold on Earth. But globally, they differ. Platinum and gold have atomic numbers 78 and 79, and their naturally occurring isotopes on Earth have atomic weights 195 and 197 respectively. In more extreme energy conditions, the elements have more isotopic occurrences. Suppose we take a

sample of matter, *a*, that is composed of atoms with 195 nucleons and 79 protons. It will fall under the extension of 'gold' given our reference scheme but under the extension of 'platinum' given the counterfactual reference scheme. The samples with atomic number 79 and the samples with atomic number 197 may be natural kinds. But it is doubtful that Locke's usage of 'gold' in 1688 or Mendeleev's in 1897 was sufficiently determinate to select either kind *uniquely*. No feature of their usage distinguishes Au 79 from Au 197 or Pt 78 from Pt 195 or uniquely determines (again in the Laplacean sense) the classification of *a* (Liston 1998).

Actual linguistic shifts challenge the claim that local usage of a word as a natural kind term can be plausibly taken to select *any* kind. The colonists, we may suppose, were using 'rabbit' as a natural kind term both before and after their voyage. But it is questionable whether the word applied to objects of any natural kind of the relevant type – the type that would secure a unique global extension for their predicate – either before or after their voyage. We cannot make sense of their pre- and post-voyage classificatory activities if we think of their uses of 'rabbit' as referring to all and only members of the same natural kind, say *Oryctolagus cuniculus* or *Lepus*. (To see this, simply replace the concepts RABBIT, HARE, and the relation of falling under a concept in displayed sentences (2) and (3) with the kinds *Oryctolagus cuniculus*, *Lepus*, and the relation of being a member of a kind, respectively, and repeat the argument on pp. 43-44 above). The only option remaining is to agree that neither their pre-voyage nor post-voyage uses of 'rabbit' referred to a natural kind. Consequently, even if natural kinds exist, the fact that a word is used as a natural kind term will do nothing to secure a unique global extension for the word. Using it as a natural kind word does not entail that it picks out a natural kind that will provide nature's contribution to the determination of global extension; it does not even make it likely that reference is globally determined, if linguistic shifts are pervasive. Appeals to natural kinds seem to be more a reflection of our tendency to project our references than an underwriting of projection.

Externalist accounts of meaning fare little better than classical accounts as responses to (PSP). Putnam correctly argues that no account can satisfy the two assumptions of classical theory. Meanings in the head cannot determine global extensions. But something must; so meanings outside the head must do so. But externalist determiners fall short of securing the kind of uniformity of meaning that would respond to (PSP).

3.

Many philosophers may find the foregoing a plausible *reductio* of the problem. They will agree that nothing we say or do, either individually or collectively, can fix the global extensions of our words in the sense of providing underlying facts connecting our words with their objective referents. However, they will argue that this is because (PSP) rests on a misconception of the relation of reference in the first place. When refer-

ence and other word-world relations are properly conceived, there is no necessity to bridge the gap between locally and globally determined references. There is no real gap to be bridged. This is what I referred to as “the dissolution response” in section 1. For purposes of this discussion, I will (somewhat grossly) divide dissolution responses into two categories: neo-Kantian and deflationist. Though deflationists are committed to none of the radical theses that neo-Kantians adopt, the two categories share a background assumption about reference that justifies discussing them in tandem. The common background assumption they share, I will argue, is controvertible and falls prey to the kind of challenge I have been presenting. Moreover, I will suggest, the particular form these positions take is to a great extent determined by the peculiar twists and turns philosophical discussions of meaning and reference took over the past fifty years. To the extent that the driving assumptions that define the rhetorical framework are not compulsory, the dissolution responses may be avoidable reactions to problems of doubtful legitimacy.

The common assumption they share is that the important part of what we know about truth and reference in our own language is independent of empirical vicissitudes. It may be a non-trivial task for us to determine to what our word ‘rabbit’ refers in the sense that we may have to do a lot of empirical investigation to discover the nature and behaviour of rabbits. However, as native English speakers, we trivially know – without any empirical investigation – that ‘rabbit’ refers to all and only rabbits just by knowing how to use the word competently. Our beliefs about rabbits could be mistaken, but not our belief that ‘rabbit’ applies to all and only rabbits. While it is a contingent truth that ‘rabbit’ refers to rabbits – language could have evolved differently – it is something like an a priori, conceptually necessary, or quasi-formal truth. In particular, we cannot coherently express in English the supposition that ‘rabbit’ does not refer to rabbits, for to do so would require us to suppose either that some rabbits are not in the extension of our word ‘rabbit’ or that that some non-rabbits are in the extension of ‘rabbit’. Either supposition is but a disquotational step away from a formal contradiction. If the extension of ‘rabbit’ is the set of rabbits, once we disquote the word ‘rabbit’, the result is that either some rabbit is a not a rabbit or some some non-rabbit is a rabbit.⁶

In section 1, I challenged this line of thought by suggesting that the English-speaking colonists circa 1620 would have been as prone to offering it as a claim about 1620s colonial English as we are about English. The falsity of

(1) ‘Rabbit’ applies to x if and only if x is a rabbit

where (1) is a sentence of 1620 colonial English, cannot be coherently expressed in our language, a sophisticated John Smith might have claimed in 1620. How could there be rabbits that are non-rabbits, or non-rabbits that are rabbits? Beneath the supposition that we could be wrong about (1), and not far beneath, lurks a formal contradiction.

⁶ See Van Fraassen 1997, pp. 90-93.

Despite John Smith's conviction, however, he is wrong about something. His confidence is only as good as his assumption that formal contradictions do not lie below the surface of the sentences he holds true, including the sentences he claims are not coherently deniable. Smith believes that no rabbit is a hare, yet he is also disposed to agree with his fellow colonists that 'rabbit' correctly applies to the local hares. Should he be made subject to the right schedule of questioning or to the right kind of interactions with the local leporidae, he will find himself led to formal contradiction. For example, if he learns that the local leporidae are neither burrowers nor altricial, properties he believes all rabbits have, he will be forced to agree that contradictions loom. His confidence that (1) cannot be coherently doubted or denied is only as good as his confidence that formal contradictions are not forced by the sentences he holds true. That latter confidence, in turn, is only as good as his confidence that his predicates have not shifted extension unbeknownst to him.

Neo-Kantians have a very different view of the matter. They agree that no word-world facts underlie our semantic projections. But they view (PSP) as a pseudo problem inspired by a mistaken metaphysical or epistemological picture of word-world relations. The origins of this picture trace back to the empiricist veil of perception doctrine, which creates a problematic gap between our words and ideas, on one hand, and an inaccessible world, on the other. This problematic gap continues to haunt contemporary philosophical discussions. The leading idea seems to be that there cannot be a word-object relation between words, on the one hand, and (mind-independent) objects that are independent of any and all conceptual structure, on the other hand, because we cannot conceptually grasp such a structure. Thus, Putnam claims that Metaphysical Realism generates a pseudo-problem:

there are to begin with all these objects in themselves, and then I get some kind of lasso over a few of these objects ..., and then I have the problem of getting my word ("rabbit") to cover not only the ones I have 'lassooed' but also the ones I can't lasso, because they are too far away in space and time, or whatever. And the 'solution' to this pseudo-problem ... - the metaphysical realist 'solution' - is to say that the word *automatically* covers not just the objects I lassoed, but also the objects which are *of the same kind* - of the same kind *in themselves*. But then the world is ... being claimed to contain Self-Identifying Objects. (Putnam 1981, p. 53)

The neo-Kantian response to this unsatisfactory state of affairs denies that word-world relations are relations between words and mind-independent objects and argues that they are instead relations between words and objects that are products of our conceptual scheming (constrained perhaps by the environment). They are relations that obtain between our words and objects as we conceive them. If we cannot reach out to mind-independent objects, the solution is to bring the objects into our linguistic and

conceptual range.⁷ Worries that our words cannot reach inaccessible extensions are quelled by making the extensions accessible. There is no lasso problem, nor can there arise any epistemological worry concerning whether our local references are globally extendible. And this is only right, because surely we know what our words mean in the trivial sense that we know the truth of sentences like (1) a priori? In this way, the principle of uniformity of meaning is upheld, not as some regularity between our words and ordinary objects, but as the truism it seems to us.

This is quite an extreme reaction to safeguard semantic projection, a reaction that is clearly not compulsory and that is irrelevant to (PSP) as I have been presenting it. We do not have to buy into the traditional metaphysical and epistemological picture that allegedly creates the unbridgeable semantic gap between our words and the objects they refer to. So, we do not have to bridge the gap by making the referents of our words be our conceptual constructions. There is a gap, but it is an ordinary, garden-variety gap. There are objects that are very far away from us in space or time, very small, very large, very energetic, etc. They are inaccessible to us in a purely ordinary sense – not in principle but *de facto*. We already know on the basis of general considerations both that such objects exist and that we and our language are evolving in ways that are not pre-established. So, we should not expect that our language is designed to deal with these objects as they become accessible to us. In other words the background for my questioning of (PSP) is ordinary realism. I assume that rabbits and gold exist and that they are mind-independent objects in the sense that their existence and much of their natural behaviour is not greatly affected by what we think (though they are by what we do). The colonists, either in pre- or in post-voyage times, are not metaphysically cut off from the New World leporidae. When they arrive, they interact in ordinary ways with the local leporidae. Nothing metaphysically mysterious or unintelligible is going on. Putnam makes it seem that they could err only if their usage had to satisfy an unintelligible constraint: their New World uses of 'rabbit' must be unintelligibly tied to the same natural kind their Old World uses were. But, as we saw in section 2, this does not seem to be a plausible diagnosis of what went wrong, independently of the existence or intelligibility of natural kinds. The plausible diagnosis is that they simply have over-generalized their word-world relationships. 'Rabbit' worked at home; it does not work in the New World – when they attempt certain practical tasks, they encounter resistance from their natural environment. None of that should give us reason to deny that word-world relations are intelligible or to adopt any variant of neo-Kantianism. (PSP) concerns not how we bridge unbridgeable gaps; it concerns only our ability to extend unproblematic local usage of predicates to objects and tasks yet to be encountered.

⁷ McDowell 1994 seems to think that philosophical puzzles about mind-world relations flow from a similar well-spring. We cannot conceptually grasp mind-independent objects, and we cannot respond with reasons to evidence understood merely as a causal, non-conceptual impingement on our senses. His solution is that we extend our conceptual range (the mind) out to reach the objects.

Deflationists about truth and reference also hold that no word-world facts underlie our semantic projections.⁸ They will be inclined to underestimate the significance of (PSP), because they proceed from the same background assumption that sentences like (1) cannot be coherently denied by English speakers. Unlike neo-Kantians, many of them are committed to no extravagant metaphysical or epistemological theses, maintaining that rabbits and gold atoms exist in the ordinary way that lay-folk and scientists take them to exist. They merely decline to grant word-world relations the kind of importance that realists about those relations grant them.

Deflationism, I believe, is also a position whose attractions are mainly due to the unattractiveness of more robust conceptions of truth and reference that developed as a result of peculiarities of debates about semantics in the past 50 years. Under the influence of Putnam and Boyd, a conception of truth and reference developed whereby those concepts were marshaled to perform a causal-explanatory role and ultimately to defend scientific realism. Against empiricists and constructivists, they argued, if the reference of 'electron' were determined classically (as whatever satisfies a concept or verification procedure or theory), then Thomson in 1898, Bohr in 1911, and Bohr in 1927 would have been talking about different objects, when each used 'electron'. Since such a conclusion flies in the face of our ordinary practices and what we ordinarily say, it provides a *reductio* of the classical determination of reference. Instead, they argued, the reference of 'electron' is determined by causal relationships between electrons and 'electron' users, an account of reference that does allow us to make sense of our ordinary practices. But the new accounts made reference assume greater and greater burdens. First, they underpinned continuity of reference. Thomson and the early and later Bohr all referred to the same kind of entity, electrons, because their usages were causally connected to electrons. This became an important tool in combating the pessimistic induction: just as most terms of scientific theories of 100 years ago do not refer (if classical theories of reference are correct), so most of our terms will not refer in 100 years time (by our successors' lights). If so, our best current theories will fail to be even approximately true. Furthermore, truth (and thereby reference) came to be employed to argue positively for scientific realism. Only if truth and reference are treated as causal-explanatory properties can we explain the increasingly instrumental successes of our successive scientific theories and the successes of the methods scientists employ in theory construction and testing. At a given stage in the evolution and development of a scientific term, it will have causally "locked on" to a referent; thereafter the referent will "guide" the selection and extension of theories about the referent on a truth-like path. On such a view, not only is chemical valence a causal-explanatory property, causally responsible for and explaining the chemical combination behaviour of atoms and molecules, the reference relation itself is a causal-explanatory property, causally responsible for our beliefs and explaining the success of our behaviour and theories.

⁸ I draw primarily on Horwich 1990 and Leeds 1978 and 1995.

Deflationists are right to reject truth and reference as causal-explanatory notions. There are good reasons to doubt that those notions are unified in the way chemical valence is, and that any account of truth and reference should be capable of supporting success of science arguments or blocking the pessimistic meta-induction. But the dichotomy that emerged from the literature – either truth is a causal-explanatory notion or ‘true’ is merely an expressive device enabling us to say things in our home language that would otherwise be inexpressible without infinite conjunction and disjunction – is unnecessarily restrictive. There is a *via media*: truth and reference can be real relations without having a unitary nature and can have all kinds of uses, including diagnostic uses, justificatory uses, and even explanatory uses.

Consider the following explanation of Smith’s behavior. Suppose Smith’s friends tell him, ‘There are lots of rabbits in the neighbourhood’, and he thereby comes to believe it. He wants leporidae for his larder, goes out to hunt some, and returns with a large catch. The success of his hunt seems to depend on his and his friends’ uses of ‘rabbit’ coordinating with each other and with the local leporidae. If they were not all successfully referring to the same things, he might well have returned home empty-handed. Thus referential connections between their uses of ‘rabbit’ and the extralinguistic environment seem to be part of the explanation of his success. Smith is successful, partly because the sentence ‘There are lots of rabbits in the neighbourhood’ is true (in the sense of being related correctly to the state of the world). The deflationist responds that we do not need to appeal to the truth of their sentence (in any sense that commits us to real word-world relations) to explain Smith’s success. Provided truth is as Tarski defined it, so that all the T-sentences are true, we can disquote: ‘There are lots of rabbits in the neighbourhood’ is true if and only there are lots of rabbits in the neighbourhood. But then the semantically deflated explanation is that Smith is successful partly because there are lots of rabbits in the neighbourhood. In this way, the deflationist contends, all the legitimate explanatory uses of robust truth and reference can be replicated by assuming no more than that all the T-sentences are true.

Of course, this is not quite right, nor does the deflationist think it is. We have been assuming that Smith and his friends’ sentences are English. But if they speak 1620s colonial English, then disquoting will produce a truth condition – that there are rabbits – that fails to explain Smith’s success. His believing that there are rabbits around and his hunting rabbits will not help him in his rabbitless environment. Deflationists think of reference as defined for the home language (where we understand the trivial reference conditions supplied by instances of (RS) simply by being competent speakers of our home language) and then as extended to other languages via translation. However, the notion of translation that is being appealed to here is problematic. How should we translate ‘rabbit’ in 1620s English into our English? We have already canvassed, in section 2, numerous possibilities – ‘rabbit’, ‘hare’, ‘harbit’, etc. – and found them all wanting, even if we rely on an overall principle of charity to constrain our translation to make as many as possible of the sentences Smith holds true come out true under the

translation. The problem is that Smith's language, in particular his word 'rabbit', does not align with the world in a way that is amenable to fitting his word-world relations to a Tarski-scheme that will make Smith rational by our lights. We have already seen that contradictions lurk among the sentences Smith holds true, and Tarski's scheme will not tolerate contradictions among sentences held true.

Similarly, the clearest diagnosis of the colonists' errors will appeal to reference or reference-related notions. The colonists' word 'rabbit' does not refer to rabbits; it refers in some contexts to hares and in other contexts to rabbits (alternatively, some tokens refer to hares; other tokens refer to rabbits). Perhaps the difference between diagnosing their error in this way is only a variant of extended disquotationalism: in some contexts, we should translate their word 'rabbit' as our word 'rabbit' and in other contexts we should translate it as our word 'hare' (or as 'European rabbit' and 'American rabbit', respectively). But intuitively, what went wrong with their language has nothing to do with how their words translate into ours and everything to do with their language aligning with the world in a way that will render them vulnerable to problems when they use the language to help them perform certain tasks. We will need to translate their sentences into, or interpret what they say relative to, our language in order to express what went wrong. But we should not think that Smith's hunting is successful or that his language is prone to breakdown because his sentences translate or fail to translate smoothly into sentences that we employ successfully. It is the real world, we should think, that constrains what he does well and does not do well with language. And this is the case also with our language. It seems to me that deflationists are inclined to think that our home language is not subject to (PSP): we can rely on its primitive reference clauses and on the T-sentences because the extensions of our predicates are determinately fixed. But this seems to be overly optimistic. Just because trouble has not shown up, it does not follow that systematic trouble will not show up when we encounter new objects and attempt new tasks, trouble that would make it necessary for us not merely to revise our ordinary beliefs but also our meta-linguistic beliefs.

4.

I conclude, as I began, with Hume. (PSP) shares several features with Hume's problem of induction. First, like Hume's problem, no elaborate metaphysical or epistemological presuppositions are required to feel the grip of the problem. Hume's discussion of induction can be motivated on the basis of ordinary commonsensical considerations and is all the more compelling for that reason. He notes that our confidence in the claim that all trees will flourish in December and January and decay in May and June is not warranted. And he is right: Cook's South Seas explorations were soon to show that 18th century confidence in that induction was based on over-generalization from a local and unrepresentative sample of evidence. In the same way, I have tried to motivate (PSP)

as a challenge to our confidence in our semantic projections, drawing only on ordinary, commonsensical considerations about finitude and contingency.

Second, Hume points out that no appeal to a principle of uniformity of nature will secure our inductions: any such appeal will inevitably beg the question against the ordinary commonsensical sceptical considerations that undermine our confidence in induction. Similarly, I believe, no appeal to a principle of uniformity of meaning will secure our semantic projections: any such appeal will inevitably beg the question. Admittedly, the colonists go wrong only by incorrectly generalizing on a spatial parameter (Old World to New World), and perhaps, now that our world is “global”, there will be little tendency on our part to commit the same kind of error. But we should not conclude from this that we will not make similar mistakes, especially given the comprehensiveness of our semantic projections: we feel confident that ‘rabbit’ will have unproblematic application not only at all places and times but even in every possible scenario. This amounts to quite a strong assumption of uniformity of meaning.

Third, Hume states,

The passion for philosophy, like that for religion, seems liable to this inconvenience, that ... it may only serve by imprudent management, to foster a predominant inclination, and push the mind, with more determined resolution, towards that side which already *draws* too much, by the bias and propensity of the natural temper. (Hume 1978, p. 40)

Hume’s point seems to be twofold. First, philosophy does not serve us well insofar as it promotes, rather than challenges, a theoretical basis for our unexamined inclinations. Second, sceptical philosophy, while it challenges these theoretical pretensions, in no way undermines our everyday reasoning or its consequences for practical activity. I share the concern expressed by the first point. Much of contemporary philosophy simply takes uniformity of meaning for granted and goes on to embed the assumption in its theories. However, I am less sanguine than is Hume about leaving everyday reasoning and its practical consequences to the authority of custom or natural habit. Everyday reasoning will be prone to breakdown if we incorrectly project.

True, the colonists need not be bothered by the fact that contradictions lurk among their beliefs until they attempt projects that unearth them. As long as they merely observe the passing leporidae, no harm is done by the fact that they call them ‘rabbits’. But if they attempt to hunt them by first finding their burrows and flushing them out, their endeavors will be frustrated, and they will return from the hunt bootless. As it begins to dawn on them that something has gone awry, they will need to make changes. In this case, the changes are easy to make. We introduce qualifiers ‘European’ and ‘American’ to filter the classifications and block the inferences that lead to inconsistency. There are two kinds of rabbits, European rabbits and American rabbits; European rabbits are burrowers and altricial; American rabbits, zoologically related to hares but not to Euro-

pean rabbits, are neither. Likewise our linguistic usage is unlikely to betray contradictions resulting from semantic over-projection until we engage in practical endeavours that make them come to light. We can act as if our language is amenable to Tarski's treatment. As soon as contradictions come to light, we will need to impose conceptual and linguistic filters that will enable us to proceed soundly with our transactions with others and with nature.

REFERENCES

- Donnellan K. (1983), "Kripke and Putnam on Natural Kind Terms", in Ginet & Shoemaker 1983, pp. 84-104
- Ginet C. & Shoemaker S. (1983), *Knowledge and Mind*, Oxford: Oxford University Press
- Goodman N. (1965), *Fact, Fiction, and Forecast*, Indianapolis: The Bobbs-Merrill Company
- Horwich P. (1990), *Truth*, Oxford: Basil Blackwell
- Hume D. (1978), *Enquiries concerning Human Understanding and concerning the Principles of Morals*, 3rd edition with notes by Nidditch, P., Oxford: Oxford University Press
- Kripke S. (1982), *Wittgenstein on Rules and Private Language*, Cambridge: Harvard University Press
- Leeds S. (1978), "Theories of Reference and Truth", *Erkenntnis*, 13, pp. 111-129
- Leeds S. (1995), "Truth, Correspondence, and Success", *Philosophical Studies*, 79, pp. 1-36
- Liston M. (1998), "Externalist Determinants of Reference", *ProtoSociology*, II, pp. 173-215
- Mencken H.L. (1963), *The American Language*, New York: Alfred A. Knopf
- McDowell J. (1994), *Mind and World*, Cambridge: Harvard University Press
- Putnam H. (1975), "The Meaning of 'Meaning'", in Putnam 1975, pp. 215-271
- Putnam H. (1975), *Mind, Language, and Reality*, Cambridge: Cambridge University Press
- Putnam H. (1981), *Reason, Truth, and History*, Cambridge: Cambridge University Press
- Van Fraassen B. (1997), "Elgin on Lewis's Putnam's Paradox", *The Journal of Philosophy*, pp. 85-93
- Wilson M. (1982), "Predicate Meets Property", *The Philosophical Review*, XCI, pp. 549-589
- Wilson M. (2006), *Wandering Significance*, Oxford: Oxford University Press

Received: July 20, 2005

Accepted: August 28, 2005

Associate Professor of Philosophy
 Department of Philosophy
 University of Wisconsin/Milwaukee, P.O. Box 413
 Milwaukee, WI 53201, USA
 Email: mnliston@uwm.edu