

2010

Are Our Words Really Real? Rabbits, Cats* and the Inscrutability of Reference

Nicholas D. Schwarzenberger
University of Puget Sound

Follow this and additional works at: http://soundideas.pugetsound.edu/summer_research

Recommended Citation

Schwarzenberger, Nicholas D., "Are Our Words Really Real? Rabbits, Cats* and the Inscrutability of Reference" (2010). *Summer Research*. Paper 41.
http://soundideas.pugetsound.edu/summer_research/41

This Presentation is brought to you for free and open access by Sound Ideas. It has been accepted for inclusion in Summer Research by an authorized administrator of Sound Ideas. For more information, please contact soundideas@pugetsound.edu.

“Are Our Words Really Real?
Rabbits, Cats* and the Inscrutability of Reference”

Today I will be discussing the theories of three philosophers who all argue that reference, and by extension our theories about the world, are underdetermined. These philosophers are W.V. Quine, Hilary Putnam, and Nelson Goodman. What the theories I will discuss emphasize is that there is an element of creation and ordering when we develop theories and models that describe the world. I say “an element of creation” because I do not want to be mistaken for advocating the full blown relativism of the sort that new age self-help books advocate, the sort that tell you can choose your reality so forth and so on. To clarify, there are two extreme over-generalizations that we tend to fall into when we talk about humans and their knowledge of the world. One is extremely passive, that we are discovering ready-made facts about the world and that we are passive learners, much as Aristotle characterized humans in De Anima. This is a view that Hilary Putnam calls metaphysical realism. The other is a view that everything is more or less arbitrary and one explanation is just as valid as another, because who are we to judge, right? The former view is much more prevalent in modern philosophy by my lights, and the philosophers whom I shall discuss were writing in opposition to this particular view. However, I do not want to be mistaken for endorsing the latter view because although these philosophers argued that there is an element of fabrication when we create theories, they definitely laid down criteria by which one could judge one theory as better than another. Anyway, to return to my thesis, each author argues that reference is not determinate and no physical or mental condition can serve to fix it; there is an arbitrariness in our terms and we are as much inventors as we are observers when develop theories. This is because the terms we use in theory are malleable and can be altered; they are not held down to an object by metaphysical glue.

Quine, Putnam and Goodman all reach this conclusion through somewhat similar examples. These examples underscore the underdetermination of our sentences and terms by demonstrating that

physical conditions are not enough to fix one interpretation to another speaker's sentence. They demonstrate that there are equally compatible and mutually exclusive interpretations, all of which would be considered a correct interpretation.

Let's turn first to Quine. Quine argues for what he calls indeterminacy of translation and ontological relativity. The example that his argument springs from is this: say that aliens come down tomorrow and a translator is put in charge of translating their language. Basically, all we will have to go off of is their behavior. Through querying them in the presence of rabbits, we determine that every time so far we'd say "rabbit" they say "gavagai;" tentatively, the extension of the two terms is the same. Does this mean we know what the word means? Quine says no, because "gavagai" could mean something like "undetached rabbit parts," "time-slices of rabbits," or "lo! rabbithood again!" Perhaps we could eventually find an analog to our grammatical apparatus of individuation (which are words like "same," the definite and indefinite articles, and copulas) and then ask "Is this the same gavagai as that?" But this assumes that the our foothold into their apparatus of individuation is somewhat fixed, and Quine suggests that perhaps the question "Is this the same rabbit as that?" could equally be construed as "Does this gavagai belong with that?" In order to determine an interpretation of a word, we must have some background or frame of reference to count as fixed. In the example above, it was the individuating apparatus, but this is not for certain. The bottom line of this example is that when we translate a speaker's utterances we have to impose some of our ontology or grammatical apparatus onto the speaker.

In order to show how Quine reaches this conclusion, I'll have to backtrack a moment and explain how it is on Quine's theory we come to master reference. Despite our intuition to say otherwise, Quine does not ascribe reference to a small child's utterances of "red" and "water" and "mama." Rather, he would say that the child is simply responding to sensory stimulus through the mechanism of conditioning. You see, Quine was a behaviorist with regard to philosophy of mind, a school of thought

that is not really in vogue anymore. Some have argued that this invalidates his argument for ontological relativity, but I won't get into that right now. If you're concerned about that, ask me at the end of the presentation. Anyway, the child is responding to sensory stimulations and in response uttering whatever word the linguistic community has taught him to utter in those circumstances, "red," "water," or "gavagai" in the case of our alien. I'm going to skip ahead a bit to the point where the child has begun to use sentences like "Fido is a dog." The word "Fido" in this sentence is a singular term, which is simply a term that refers to one object; "dog" is a general term that Quine defines as "true of any number of objects, from none up" (Roots, pg 84). This sentence is an example of predication, where you're essentially saying " x is F " in some way. Reference develops through the use of the relative pronoun, which allows a child to rephrase sentences involving active verbs (for example, "I bought Fido from the man that found him") into an example of a predication ("Fido is a thing x such that I bought x from man y such that y found x " to use a somewhat convoluted example). Bear in mind "such" is not the first thing that comes to mind as an example of a relative pronoun, most of the time we think of "who," "which," or "what." But it's just rearranging deck furniture on the Titanic, we could rephrase the sentence "Fido is a thing which I bought from the man who found him;" the earlier example was rephrased with an eye to formal logic. Regardless, what matters here is that the relative pronoun allows us to turn a clause in a sentence into something that can be used in the same way as general term in a predication. An example of this use of the relative pronoun in a predication is "Everything that we salvaged from the wreck is in the shed," which is an instance of the universal categorical ("all F are G "). This is significant because the child has now begun to use quantification, which for Quine is essential to reference.

At this point Quine lists some derivations from which you can get from the universal categorical to the other ones. What bears noting that there are two kinds of quantification. There is substitutional quantification and objectual quantification: substitutional quantification operates under the assumption

that variables are placeholders for the names of specifiable objects, with universal categoricals counting as true

...if and only if the open sentence following the quantifier comes out true under every substitution for the variable; and an existential quantification counts as true if and only if the open sentences comes out true under some substitution.

(Roots, pg 98).

Basically what this is saying is that it limits discourse simply to objects whose names we can plug into the variable. Objectual quantification differs in that “the variable refers to objects of some sort as its values; and these need not even be objects each of which is separately specifiable by name or description” (Roots, pg 98). So basically objectual quantification allows us to make statements about unnamed or unspecifiable objects, which greatly expands the number of things we're able to talk about in big ways, allowing us to discuss grains of sand, dogs we haven't encountered, and electrons, to mention a few examples. Initially the child's quantification will be substitutional, but according to Quine we make the jump from substitutional to objectual quantification because for

...such examples as 'An apple is a fruit', [and] 'A rabbit is an animal', it would be inappropriate to read ' $(x)(\text{if } Fx \text{ then } Gx)$ ' in the substitutional way as meaning merely that every substituted name that verifies ' Fx ' verifies ' Gx '. It is unnatural if not absurd to imagine names, or singular descriptions either, for all apples and rabbits. (Roots, pg 99)

Objectual quantification is where the language speaker begins to use reference to the most of its capabilities. Quine puts it as “Once a theory is formulated in quantificational style, its objects of reference can be said simply to be the values of its quantified variables” (Roots, pg 100). Objectual quantification completes a full switch to a Tarski-esque way of thinking about language; “What does the general term “rabbits” refer to? “Rabbits of course!” It's a more nominalistic way of approaching language, where it's the use of the word and set of phenomena we refer to. I want to emphasize that this

is an ontologically neutral way of approaching language.

How is this ontologically neutral and what does that have to do with inscrutability of reference? In our language we can ask a speaker, if their use of the word “rabbit” should seem funny to us, what they mean by saying “rabbit.” And they can answer, if they are using it in a different way, that they're referring to “undetached rabbit parts” or “time-slices of rabbits,” what have you. However, this is just allaying a bit of inscrutability of reference by making use of other words, which in turn suffers from the same problem. Quine states that in the case of our native languages we're often content to settle on the principle of homophony, which is to say if we're saying the same thing we must mean the same thing, and we then temper that with the principle of charity, which states that if we have to alter how we interpret a speaker's words to make them make more sense, we will. We often make use of the principle of charity when we have to adjust for strange uses of words, like “sick,” “gnarly,” or “chill” all expressing approval of something and not what they traditionally mean. Anyway, objectual quantification allows for such inscrutability and indeterminacy because it is essentially what makes the predicates of quantification somewhat free-floating; when we say “everything x is such that” or “there exists x such that,” we are using objectual quantification and it makes the quantification true of whatever makes it true. “Every” and “some” become the staples of our ontology, and everything else is specifiable within the system. Ontologies then become relative, with homophony and our everyday ontology (whatever that is) being our initial starting point; we could then switch to some manner of “time-slice” ontology and the people we interact with would still be understandable and we would be most intelligible to them, unless the issue of “same” came up. But once again, appealing to words like “same” and the individuating apparatus is just relying on homophony, relying on the assumption that our word “same” means the same thing. The objectual quantification makes it so that theories simply refer to whatever their terms are true of; we can use various definitions of a term or switch between ontologies, even drastically different ones. For example, so long as all of the logical relationships and

truth values are preserved, we could make a theory of zoology refer to Gödel numbers. To repeat, so long as true sentences remain true and logical connectives are preserved, we can specify whatever universe of discourse we please for theory.

And so the force of Quine's theory then, when we turn to how human beings manufacture theories about the world, is that our terms are often defined in relation to other terms. Quine argues for holism with regards to theory, that sentences in a theory derive their meaning and force from their relationships with one another. He distinguishes between observation sentences and theoretical sentences; the former, because they rely on sensory stimulation, are thus are susceptible to the same issue our translator had when he was attempting to determine what “gavagai” meant. Quine's argument is that if observation sentences, in which we have the advantage of observable circumstance, suffer this problem, then it infects our theoretical sentences more deeply; because in those cases we rely on a sort of indirect ostension in order to show how they relate to our sensory stimuli.

What are the practical consequences of this? When we develop theories about the world, we can switch between ontologies and universe of discourse; I would argue this has been something somewhat implicit in theories hitherto, when we opt to define categories like “human” in different ways, whether it be a collection of space-time points, a set of atoms, or a reasoning animal. It further supports Quine's view that scientific theories have their fronts in observation sentences and predictions, whereas theoretical sentences just connect observation terms. It also leads to the radical conclusion that another speaker could be talking about Gödel numbers or geometry, and if their utterances match with our expectations we would never have any indication. For Quine reference itself is indeterminate.

From here let's move on to Hilary Putnam. He argues that objects and our terms do not exist independent of a conceptual scheme; he sets himself in opposition to a view that he terms “metaphysical realism.” Metaphysical realism entails that reference is a definite relation between a word and an object, and that the world is filled with ready-made facts waiting for us to find them.

Putnam sets himself in opposition to a common view that operational and theoretical constraints on scientific theory allow us to pick out true sentences and by extension the extensions of our terms. Let me briefly elaborate on the view that Putnam opposes.

According to some philosophers of science, science operates by creating operational constraints. An operational constraint is a sentence along the lines of “an admissible interpretation is such that most of the time the sentence S is true when the experiential condition E is fulfilled.” The term experiential condition refers to experiments or observations, and so it basically says “we can tentatively conclude that sentence S is true if an experimental result is thus and so, leaving room for revision as need be in light of future findings.” A classic example of this is the sentence “A current is moving through the wire” being true if the needle of an instrument connected to the wire moves. It's a simply a specification of a way that scientific method operates. What theories we ultimately accept as true are also subject to what are called theoretical constraints, which are essentially rules like “accept sentences that require the least revision in already existing sentences without undue complications,” which is crude formulation of the rule of conservatism. Operational and theoretical constraints on this view allow us in theory to determine true sentences and the extensions of terms, as well as the intensions of terms (spelled with an “s”). An intension is simply a function that specifies what the extension of a given term would be in every possible world. If we knew what phenomena we could consider mammals in all possible worlds and circumstances, then that helps us greatly to understand the term “mammal.” This relates to an operational constraint in that the knowledge of an operational constraint in theory allows us to grasp an intension of the terms of a sentence, and hypothetically by extension through experimentation the extensions of the terms as well. The operational and theoretical constraints allow us to determine which sentences in a theory are true; however, the view that they would allow us to determine the extension and intension (namely, the reference) of terms isn't feasible according to Putnam.

Putnam argues that “no view which only fixes the truth-values of whole sentences can fix reference, even if it specifies truth-values for sentences in every possible world” (Reason, 33). The typical view assumes that operational and theoretical constraints, because they determine the truth-values for whole sentences, will determine the reference (intension and extension) of the terms. Putnam's counter-example is as follows. Take the sentence “A cat is on a mat.” Let's assume “is on” is tenseless, so it means “is, was, or will be on.” Typically the sentence “A cat is on a mat” is true in all worlds in which a cat is on a mat and is false in all worlds in which no cat is on any mat. However, we can redefine the intension of the term as follows: we shall distinguish between three possible cases for our intension, case (a) where a cat is on a mat and a cherry is on a tree, case (b) where a cat is on a mat and no cherry is on a tree, and finally case (c) neither of those. We then define x as a cat* if and only if case (a) holds and x is a cherry; or if and only if case (b) holds and x is a cat; or finally if and only if case (c) holds x is a cherry. Mat is defined on similar lines as x is a mat* if and only if (a) holds and x is a tree; if and only if (b) holds and x is a mat; and finally if and only if (c) holds and x is a quark.

Basically, the idea is in world (a) where a cat's on a mat and a cherry's on a tree the word cat* will refer to cherries and mat* will refer to trees; in world (b) cat* refers to cats and mat* refers to mat; and finally in world (c) cat* refers to cherries and mat* refers to quark. If you look at the diagram on the sheet of paper I gave you you'll see that in all worlds where the sentence “a cat is on a mat” is true the sentence “a cat* is on a mat*” is true as well and all worlds in which one is false the other is false, too. The first thing that should strike you is that these predicates, cat* and mat*, are contrived. Painfully contrived. However, this sentence adheres to the operational constraints (we have a means of telling if a cat* is on a mat*; namely, we look) and we could go about making a theory to fit these words in. More importantly, the intension I've specified isn't obvious at all when you look at the pictures. My pictures have all the relevant information, and if I hadn't been specifying the intension or saying the “asterisk” part aloud, you would've thought that I was simply using our normal run of the

mill cat and mat predicates. If we had been describing this picture and I were using those predicates without telling you, we would have agreed on truth values and everything. This begs the question, what makes our words not have the reference of words like cat* and mat*?

Operational and theoretical constraints allow us to know the truth values of various sentences within our theories; and this is problematic because as we saw above the truth values of full sentences underdetermine reference. Putnam argues that all the sentences in a theory can be interpreted in ways analogous to cat* and mat* in his book and still have the truth values preserved. This is much analogous to Quine's ontological relativity I discussed earlier, that what our predicates in our theories refer to is malleable, and moreover that what they refer to can be shifted so long as truth value and logical relations are preserved. Making cat* refer to cherries and not cats is rather small compared to shifting a theory to talk about Gödel numbers instead of atoms.

The larger point that Putnam is driving at is that objects and the words for them do not exist apart from a conceptual scheme. If sentences underdetermine reference, what does determine it? Nothing. A word means nothing apart from its use in theory. This doesn't mean that there's no experiential inputs to our theories, simply that a word does not have a necessary link to an object. Furthermore, it makes no sense to talk of an object outside of a conceptual scheme; an object or word is picked out and distinguished by various significant properties out of a possibly infinite set, but how do we do that? We specify them, which also makes use of reference, and thus we begin to find ourselves in a regress. A word is meaningful in relation to other words, but this doesn't mean experiential elements are blocked out: we use words and reference to sort out sensory experience as well. However, as we saw with the Quine's example, that suffers from some slight indeterminacy as well, but we could say we're reasonably confident in the case of the correspondences between words and objects we interact with or point to.

A common objection to Putnam is that a causal chain of the appropriate sort can help us specify

what reference is. The problem with this lies in the phrasing “appropriate type;” at that point, we're making us of reference to specify reference, which to say the least begs the question. Why does it beg the question? It begs the question because you're making use of the device you're trying to define to define the device; furthermore, the term “reference” is going to be related to other terms and sentences in the theory of reference, and these are all equally underdetermined.

This attempt to root reference in the actual world is an attempt to epistemologically exceed our reach; our words mean nothing outside of a conceptual scheme. Since our words are undetermined, there's an element of fabrication when we specify terms and create of theories. This is what Putnam calls internal realism, which contrasts what he calls metaphysical realism. Metaphysical realism, again, is the belief that the world consists of ready-made facts which our theories seize on by some causal reference relation or some analogous device that serves the same purpose. The favorite mode of viewing for this view is the God's eye view; where we can look down on a system or theory, passively, in its completeness. However, we can't make the epistemological hurdle to get to the God's eye view; we're constrained to our point of view, and our use of terms to talk about the world. Moreover, Putnam asks specifically what makes reference determinate: what is it that makes “cats” refer to cats and not cats*? Unless we are able to answer what “fact” makes reference determinate, it seems as though the God's Eye view is untenable, regardless of whether or not we can adopt it as humans. When we are talking about the world, our theories do account for and predict the behavior of something that stimulates our sensory nerves. Internal realism allows for a sort of relativism, but it is tempered by theoretical concerns. On this view, an ideal theory cannot be false.

To conclude, truth values do not guarantee us our reference, namely our intension and extension. A speaker could use vastly different, irregular categories or shift the reference of predicates drastically, and we would not know unless it made a speaker's behavior irregular enough to question it. But when we question it, they will answer with words and we must either assume those words have the

same reference and function as ours, otherwise they're unintelligible. Take for example a lady who uses our cat* predicate. She possibly takes her cat* to the vet, or likes to put cats* atop her Shirley Temples. If she shows up to the vet with a cherry, that's rather distressing towards our view of her mental health, but on the other hand we can figure out if every time she talks about a cat* she means a cherry. But that only gives us the extension in this world; we would never have any indication that if the world had no cherries on trees that her word cat* would differ any from ours. The observed conditions do not account for what our terms refer to.

Goodman argues that when we create a theory or work of art we are creating a world or world-version. On his account in science we are engaged in as much an act of creation as an act of discovery. However, Goodman points out through citing psychological studies that simply by perceiving we are also fabricating to a degree. Similarly to how Putnam argued that the world does not contain ready-made facts, Goodman argues that we impose a structure on the world when we perceive.

Goodman works from studies about real versus apparent motion. The short version of it is this; psychologists flashed one dot against a contrasting background and then another dot. With a less than ten millisecond gap between the two, the viewers on average saw them as simultaneous. At more than 45 milliseconds, observers saw two discrete dots, one flashed after another. However, between ten and 45 millisecond they saw one dot moving from one position to the other. Further tests were done, and it was discovered if the dots were of different shapes the dots would smoothly transition into one another and so forth. Goodman devotes a whole chapter to this in the book I read, Ways of Worldmaking, if this sort of thing interests you.

One of the more telling examples in Goodman's book concerns apparent motion and engineers. The results concerning apparent motion are simply averages or means; there's at least a degree of difference in how different people perceive apparent motion. Some people, typically engineers, are unable to see apparent motion at all, and in the study these were classified as naïve realists (just a term,

it's not stating anything concerning their philosophical disposition). Goodman makes an argument similar to the rabbit and cat* examples, but I'm going to have to rephrase it to make the parallel more obvious. Say we have a participant in the experiment concerning apparent motion. He gives us the verbal report "I see two distinct flashes." What does he mean by "see," Goodman asks. Goodman speculates "perhaps he means that he sees the two as we might say we see a swarm of molecules when we look at a chair, or as we do say we see around table top even when we look at it from an oblique angle" (Ways, 92). The study indicated that an observer can begin to distinguish between apparent motion and observed motion eventually, so the possibility arises then that he sees them as one distinct object moving but recognizes the qualities that distinguish apparent from observed motion, and answers that he sees two distinct flashes.

However, for the experiment this creates problems. Goodman again asks "Shall we say, rather, that he misunderstands the instruction, which is presumably just to tell what he sees? Then how, without prejudicing the outcome, can we so reframe that instruction as to prevent such a 'misunderstanding'?" (ibid., 92). Presumably he is telling us what he sees, but the word "sees" is what presents us with the problem. The solution Goodman proposes is that we instruct him to tell us what he sees in phenomenal terms, but this makes the distinction between real and apparent moot. He'll be perceiving either that one dot moves or two dots flash. We allay our suspicion of the word "sees" by appealing to another word or conceptualization. This conceptualization is part of a prior frame of reference or conceptual scheme; it's not regarding some external, ready-made fact of the world. That we resort to phenomenal terms, a pre-existing scheme, is similar to the manner in which we impose our ontology onto those we translate according to Quine. And you'll note also that our appeal to phenomenal terms also relies on homophony; we're just assuming that the person we communicate with knows what we mean. If they don't, we simply hunt for a definition of "phenomenal terms" in terms they do understand, or say they understand anyway.

This brings me to the question of what larger point these examples illustrate. They have structural similarities: to begin with there's an utterance in which we have no pre-established framework to interpret it (shared ontology, conceptual scheme, what have you). From there unless we situate it within a framework, it's underdetermined. We can render or translate the utterance satisfactorily in a variety of ways. When we decide to situate it in a framework we assume that we understand the terms of the framework; however, these are just as susceptible to questioning as the term we're trying to fit in. It seems we're caught in a regress; if we try to specify a term, we must do it in terms of other words that we assume are specified, but can only be specified in turn by more terms. This is at the heart of ontological relativity, Putnam's internal realism, and Goodman's conception of worldmaking. New conceptual schemes (I'll use this term hereon out for frames of reference) are generated from pre-existing ones. They're defined in terms of the old ones; new terms are defined with old terms, changes in terms are tracked by other terms, and perhaps some old terms are combined into one new term. The experiential inputs of theories are not enough information to determine reference; for when it comes down to it, time-slices of rabbits and rabbits don't even have the same extension: there are many more time-slices of rabbits in the world than there are rabbits, because each rabbit consists of multiple time slices. So as Quine points out, this indeterminacy permeates extension and intension.

One objection that has been raised is that there is more information than "experiential inputs," and this information allows us to determine reference. Perhaps in addition to facts like truth-conditions, ostension by a speaker and other such evidence in their behavior, we can appeal to brain states to determine what words refer to. Perhaps there's a physical state that corresponds to certain uses of a word, certain intensions or extensions. The problem we have here is that the inscrutability of reference is still in effect; how would we test brain states to find a one to one correlation with words? If we look at the brain of someone observing a rabbit in the field and uttering the word, it's not just the portion of

their brain corresponding to the utterance that will light up. There are also ostensibly neural processes involved with perception. However, we have no idea which ones are involved with the perception of the rabbit, because they aren't simply seeing a rabbit, there's also a background against which the rabbit is placed. Perhaps we then put them in a dark room where the only thing that will flash is the image of the rabbit and we have instructed them to name what they see. They utter the word "rabbit." The brain scan lights up. However, they're still perceiving the black background. But perhaps if we changed colors but not the image, we can determine which neurons correspond to the background? But how do we account for the fact that different rabbits look different? A brown rabbit will stimulate different neurons than a white rabbit. And what about size? The position of the rabbit? But perhaps I'm digressing. We're trying to find the portion of the brain that corresponds to the "rabbit" word. Well there's the vocal apparatus, but that's muscular stimulation. The problem comes down to how do we distinguish the neurons that correlate with the portion of the brain that decides which characteristics of rabbithood are significant for an animal to be a rabbit? To those who've studied philosophy of language, does this sound familiar? We've swapped meanings for brain states. We've fallen into talk of meanings again. Language is use, and circumstances of use vary always. The compositionality of English alone allows us to create a stunning number of new sentences, and linguists have calculated the majority of sentences we utter in a day are new. What's most interesting is that the problem we ran into with trying to understand the speaker's sentences arises in their brain. We're given a complete set of brain states and we're attempting to determine which configurations of neurons firing correspond to a "term," the portion of the brain responsible for linking perception with linguistic utterance. The same process of induction we were using to translate physical utterances is brought to bear with all of its indeterminacy on a neuro-chemical language.

Clark Glymour argues that the notion of a physical relation for reference is still viable; he observes that Putnam argues that trying to use reference to specify reference leads us in an endless

circle. Glymour states that while Putnam thinks this circle is vicious, he does not. Glymour states that “It does not follow that causal relations fail to fix the word “causes”, only that sentences containing the word “causes” fail to fix the reference of “causes”” (On Conceptual Scheming, 177). Glymour implies that causal relations can still serve to fix words and interpretations. My instinct is to ask how it is we can establish what the causal relations are without talking about them. The evidence for these causal relations will arguably be underdetermined; if the world is full of ready-made facts as Glymour and the metaphysical realists say, how will we determine which model of causal relations is the correct one? Glymour acknowledges that it is “clearly utopian to suppose that one might *define* “refers to” or “signifies” in causal or physical terms. At most one should expect *natural explanations* of aspects of reference, or perhaps useful constraints, in like terms, on reference or co-reference” (ibid, 179). But what are these “natural explanations” and “useful constraints?” It sounds like he's saying the most we can expect is a rough model of reference, which perhaps Putnam might agree with, but the implication is that these causal relations that we can't adequately describe will hold and fix interpretation of our other terms. If we don't know the relations that serve this function though, how shall we fix our interpretation? It seems to me that the debate between internal realism and metaphysical realism comes down to whether or not there's a fully formed real world. The internal realist argues that for our intents and purposes, no, because we can't speak of it and it's superfluous to us. As part of the world, we're limited in how much we can discuss about it and ourselves. The metaphysical realist instead states that our theories are talking about something, and while they may be underdetermined, what they describe really does exist, even if we get it wrong sometimes. For the metaphysical realist, even an ideal theory would possibly fail to describe the world; how we would know though? Perhaps we would find a better theory, but that can't be right because this is the ideal theory. The way the world is would really be lost to us. If our ideal theory can still turn out false, what does it mean for something to be true? If our theories are unable to speak about or account for this world of definite facts, what reason do we have to

believe it exists?

To conclude, I think the most important characteristic of the accounts given by Quine, Putnam and Goodman is that they allow for the invention inherent in linguistic use. A few weeks ago I saw my first opossum in my friend's backyard at three in the morning. No one else was around, and I had never seen a opossum before. I'd see animated ones on TV, cute and all, but this looked nothing like them to me. Of course now I see the resemblance, but at the time the only explanation I could come up with to explain what this big, ugly rodent eating my friend's trash was that a rat had found toxic waste and mutated to (comparatively) giant proportions. I called it a giant rat, but then described it to my friend and he told me that was a opossum. The moral is when we encounter new phenomena, we have to make our existing terms do work to incorporate these new pieces of sense data. We invent new words or make old words do new tricks; our new world is defined in the terms our old world uses. Goodman states that different worlds or versions of the world (to him there's little to no difference) can be useful for different purposes, and that frame of reference is important. Our knowledge of the world is always in relation to the frame of reference we occupy, which in turn allows for different frames of reference for different purposes. For tracking celestial objects, it's useful to assume the earth is fixed, but when modeling the rotation of galaxies, perhaps the center of a certain galaxy is better than the earth for a reference point. We see this model relativism in physics, where physicists switch between conceptualizing of gravity as field and thinking of it instead as spacetime curvature, simply as a matter of ease. Is there a fact of the matter as to which model is right in any of these cases? Goodman, and I as well, would argue no. It's a matter of theoretical concerns, such as how a theory accounts for the phenomena, where it stands in regards to already existent theories, and how well it makes predictions. I end with one of my favorite Goodman passages, which I think encapsulates this issue best and expresses an irreverence I find charming:

Is the screen that a dot moves across the same as the one no dot moves across? Is

the seen table the same as the mess of molecules? To such questions, discussed at length in the philosophical literature, I suspect the answer is a firm *yes* and a firm *no*. The realist will resist the conclusion that there is no world; the idealist will resist the conclusion that all conflicting versions describe different worlds. As for me, I find these views equally delightful and equally deplorable—for after all, the difference between them is purely conventional! (Ways of Worldmaking,119)

Bibliography

- Glymour, Clark. 1982. "Conceptual Scheming, or, Confessions of a Metaphysical Realist". *Synthese* 51: 169-180.
- Goodman, Nelson. 1978. *Ways of Worldmaking*. Indianapolis - Cambridge: Hackett Publishing Company.
- Putnam, Hilary. 1981. *Reason, Truth, and History*. Cambridge: Cambridge University Press.
- Quine, W. V. 1973. *The Roots of Reference: The Paul Carus lectures*. La Salle, Illinois: Open Court Publishing Company.
- . 1969. "Speaking of Objects". In *Ontological Relativity and Other Essays*, 1-25. New York: Columbia University Press.