

---

PERCEPTION,  
REPRESENTATION,  
LANGUAGE

---

FRANK JACKSON


---

THE AMHERST LECTURE IN PHILOSOPHY

LECTURE 7, 2012

---

*<http://www.amherstlecture.org/>*



*Perception, Representation, Language*

Frank Jackson

PREFERRED CITATION

Jackson, Frank. "Perception, Representation, Language." *The Amherst Lecture in Philosophy* 7 (2012): 1–17. <<http://www.amherstlecture.org/jackson2012/>>.


ABSTRACT

We use sentences to make claims about how things are. Many of these claims come, in one way or another, from how our perceptual experiences represent things to be. I argue in the lecture that this tells us that very often we need sets of centered worlds (instead of sets of worlds) to capture the contents of our sentences. I review some of the implications of acknowledging this for the debate over proper names and the referential behavior of the word "water."

*The Amherst Lecture in Philosophy* (ISSN: 1559-7199) is a free on-line journal, published by the Department of Philosophy, Amherst College, Amherst, MA 01002. Phone: (413) 542-5805. E-mail: [alp@amherst.edu](mailto:alp@amherst.edu). Website: <http://www.amherstlecture.org/>.

---

Copyright Frank Jackson. This article may be copied without the copyright owner's permission only if the copy is used for educational, not-for-profit purposes. For all other purposes, the copyright owner's permission is required. In all cases, both the author and *The Amherst Lecture in Philosophy* must be acknowledged in the copy.



# Perception, Representation, Language

Frank Jackson

Princeton University

The Australian National University

La Trobe University

WE LANGUAGE USERS HAVE VIEWS about how things are and, on occasion, want to share them with our fellow language users. That's why we produce sentences like "Global warming is a myth," "Watch out, that plate is hot," and "Arsenic is poisonous." That's why we write histories and publish physics textbooks – and why I am giving this lecture. That is to say, a principal use for words and sentences is to give information, to make claims about how things are, to represent the world and the things in it as being a certain way. Here we have three different phrasings designed to get across the same basic idea – you can say it in terms of information, or in terms of representation, or in terms of making claims.

How could a word or sentence – a certain physical configuration on a page, or on a computer screen, or encoded in a sound wave – give information about the way things are? In the broad, we know how the answer has to go. There are functions that go from the physical structures in question to ways things might be, functions we grasp as competent users of the language in question, and we exploit our knowledge of the functions to extract the information, to find out how things are being represented to be.<sup>1</sup> Thus, there is a function that goes from the sentence "X is circular" to the shape X has to have in order to be circular. Compe-

---

<sup>1</sup> Of course, competence comes in degrees, but we will set this complication aside in what follows.

tent English speakers grasp that function and this is how, when they hear “X is a circular” uttered assertorically, they know what the speaker is saying about X’s shape – it is the right way for “circular” to apply to X. And they know what that is – it is to be circular.<sup>2</sup> Again, there is a function that goes from “Global warming is a myth” to a way the world might be, namely, a world where global warming is a myth. English speakers grasp that function and this is why, when they wish to claim that global warming is a myth, they produce “Global warming is a myth” to one or another audience.


We can say all this in terms of reference, but here I think we have to be a little careful. The term “reference” gets used in (at least) two different though related ways in discussions in the philosophy of language, and sometimes they get mixed up and confusion results.

Suppose I ask you “Do you know the reference of the phrase ‘The top of the stock market?’” You might say, “Of course not. That’s a very valuable piece of information. It tells one when to buy and sell. If I knew the reference of the ‘top of the stock market,’ I would be holidaying in some luxury resort paid out of the profits of my trades, not having a conversation with you here in Amherst.” But of course there is a sense in which you do know the reference of “the top of the stock market.” You know what it takes to be the top of the stock market. You know how something has to be to be the top of the stock market. What it takes is for it to be the *peak*, to put it in graphical terms, and you know that without knowing where the peak comes. Indeed, your knowledge of the reference of the “the top of the stock market” in this second sense explains how you know that knowing its reference in the first sense would be such a valuable piece of information.

Our concern in this lecture will be with the second sense of reference: the “what it takes” sense, the “how something would have to be” sense. This is the sense that feeds most directly into any account of how words and sentences give putative information, and of what a speaker is claiming about how things are when they use a word or a sentence. For this is the sense which connects the words and sentences that come from our mouths, pens and computer keyboards, with how things are. This is the sense which connects with our earlier talk

---

2 Philosophers disagree about the conditions under which someone knows what it is to be circular – Does one need to be able to give the geometers’ definition? Must one be able to recognize a circle?... – but that, I insist, despite dissent in some quarters, is no reason for doubting that pretty much anyone reading this lecture knows what it is to be circular.



of functions. The functions that underpin the way we extract information, know what's being claimed, know how things are being represented to be, are reference relations understood in the "what it takes" sense of reference.

I hope the foregoing strikes you as common sense about words and sentences, dressed up in a bit of jargon, or, if you like, as folk theory spelt out a bit. Do any of us need to do a course at university to know that sentences give information, and to know what information a sentence like "I am hungry" imparts to English speakers? However, despite its commonsensical flavor, it raises some important and tricky questions. Here's an example.


### The Puzzle about the Word "Water"

PEOPLE USED THE WORD "WATER" in the early 1700s to give information, as in sentences like "There's water over the horizon" and "Plants need water to survive." However, they did not know about  $H_2O$  in the early 1700s; what then was the information they were giving? Plausibly, roughly, that there's watery stuff over the horizon, and that plants need the watery stuff to survive. That's what they knew. They knew, that is, that there is a wide-spread kind that falls from the sky, fills rivers, is a liquid at room temperature, is potable, colorless, etc., and that plants need it, and sometimes they knew that there is some of it over the horizon. How then can we explain the fact that "water" fails to refer to XYZ on Twin Earth?<sup>3</sup> For XYZ is the watery stuff on Twin Earth. Why isn't XYZ the right way to be stuff that the word "water" applies to?

Some will say that there's no puzzle. There is only a puzzle if you assume conventional wisdom about the Twin Earth thought experiment, and that's a mistake. I half agree with this response. People might have used "water" in such a way that it applies to XYZ on Twin Earth. However, I think it is also clear that we might have used the word "water" in such a way that it doesn't apply to XYZ, and that many philosophers in fact use it in this way. So there's still a puzzle, the puzzle of what to say for *this* usage of "water," and from now on I will use the word

---

3 I am of course referring to Hilary Putnam's famous Twin Earth thought experiment in "The Meaning of 'Meaning,'" in *Mind, Language and Reality* (Cambridge: Cambridge University Press, 1975), 215–271.



“water” in the way that means that XYZ isn’t the right way for “water” to refer to it. What that way is exactly, we discuss later; for now, what is important is that there is such a way.

Others will object to a presumption in the penultimate paragraph. I presumed that, as a rule, people in the early 1700s were justified in much of what they claimed about how things are when they used the word “water,” and that any story about the reference of the word “water” (understood in the sense of “reference” tied to information) should respect this fact. Maybe what some in the early 1700s claimed about how things are when they used the word “witch” wasn’t justified, but surely the situation was very different for the word “water.” This presumption underlay the earlier claim that we should not read their assertions as ones about  $H_2O$ . For in the early 1700s, people weren’t justified in believing, for example, that plants need  $H_2O$  to survive, or that there is  $H_2O$  over the horizon. Some may object (have objected in fact) that in making this presumption I confuse epistemology with philosophy of language.

I agree that epistemology and philosophy of language are different subjects in a department’s list of offerings, but that doesn’t mean that what we say in the philosophy of language class isn’t answerable to epistemic constraints. We use words and sentences to express how we take things to be, as we said right at the beginning. I, along with thousands, am convinced that Plato was a great philosopher, and that my conviction is fully justified: it is certainly true and something I know about the world I inhabit. I use words and sentences to express this conviction, most obviously the words “Plato was a great philosopher” (the words that in fact occur in the preceding sentence). This means that what I claim about how things are using the sentence – a question about the referential properties of my words – has to be something that is certainly true and known by me to be true. The same goes, I insist, for the use of the word “water” in the early 1700s in sentences like “Plants need water to survive.” Any account of the reference of the word needs to be consistent with these sentences making claims about how things are that were known to be true in the early 1700s.

Still others will object that people in the 1700s *did* know that water is  $H_2O$ , and were justified in believing that water is  $H_2O$ .<sup>4</sup> What in that case, I reply, were those experiments in the late 1700s and early 1800s all about?

---

4 In phrasing matters this way, I follow the common practice of supposing that water is  $H_2O$ . Maybe, strictly speaking, we should say that water is composed of  $H_2O$ , or is composed of hydrogen and oxygen in a ratio

So there really is a puzzle. I will tell you what I think the solution is about three quarters of the way through this lecture.

## How Should We Think about Information?


WE HAVE BEEN TALKING OF words and sentences as delivering putative information, as serving to make claims, as representing how things are. A natural way to spell out this talk of information, claims and representation is set-theoretically using possible worlds. We can think of a possible world as a complete way things might be, exactly one of which is the actual world (the one we occupy). Some of these complete ways things might be will be in accord with how a sentence like “Global warming is a myth” represents things to be, and some will not. The former will be the possible worlds where global warming is a myth, and the latter will be those where it isn’t. We can then say that the way the sentence represents things to be, the putative information it delivers, its representational content,<sup>5</sup> is given by the set of worlds each of which is as “Global warming is a myth” represents things to be. Speaking more generally, the (representational) content or putative information delivered by a sentence is given by the set of worlds, each of which is as the sentence represents things to be. And we can think of the sentence as saying that each world in that set, and that set alone, might be the actual world – or, better, we can think of a person who uses the sentence to make a claim about how things are as saying that each world in that set, and that set alone, might be the actual world.

This way of thinking has a great deal of intuitive appeal. Don’t we go to lectures and read books in order to garner opinions, and hopefully sometimes knowledge, about the kind of world we live in? Is it one where climate change is happening; is it one where there will be peace in our time; is it one where the Higgs boson exists? But information about the kind of world we live in is information that classifies or categorizes our world, just as information about the kind of car you drive is information that classifies your car as, say, a sedan rather

---

of 2 to 1, or that water molecules are H<sub>2</sub>O molecules, or . . .

5 Its content, for short, in much of what follows. To highlight the fact that there are other things one might mean by “content,” I elsewhere call content in this sense *ir-content*. See Frank Jackson, *Language, Names, and Information* (Chichester: Wiley-Blackwell, 2010).



than an SUV. And to classify or categorize is to make a partition. In the case of information about our world, it will be a partition into the worlds which fall under the classification versus those that don't. In the case of kinds of cars, it will be a partition that divides the sedans from, say, the SUVs and the pickups.

Of course thinking of information in terms of partitions among possible worlds poses many questions. How do the possible worlds it trades in relate to the distinction between what's conceptually possible versus what's metaphysically possible? What's the metaphysical status of the non-actual possible worlds? How does the set of worlds thought of as the informational or (representational) content of some given sentence relate to the sentence's meaning? These are challenging questions, but I want to insist, all the same, that there is something basically correct about the set-theoretic/partition among possibilities way of thinking about information and representation. Someone worried about whether a tumor is malignant or benign wants to know which of two possibilities is actual.

The main message I want to get across in this lecture, however, is that, in very many cases, we should replace the talk of partitions among possible worlds with talk of partitions among *centered possible worlds*, and that doing this has significant repercussions for how we get information from coming across sentences we understand. It also gives us a simple way of resolving the puzzle we raised earlier about the word "water," and in addition gives us an attractive way of thinking about the informational role of proper names.

We will start with the need to traffic in centered worlds.

## Centered Worlds and the Content of Perception

THE VIEW THAT WE NEED CENTERED WORLDS has by now quite a history. Much of it involves reflections on the phenomenon of egocentric belief (and knowledge) and on the sentences that report a person's egocentric beliefs. There are nice arguments that show that what I believe when I believe that I myself am thus and so cannot be expressed as the belief that the so and so is thus and so. They are irreducibly beliefs about myself, not about myself *qua* the thing that is so and so. A number of philosophers have played important roles in convinc-



ing us of this, perhaps the best known players being John Perry and David Lewis.<sup>6</sup> I accept the conclusion of these arguments but I think the focus on them has been in some ways unfortunate. It has led some to think that we only need centered worlds for rather special cases, whereas in fact the need for them is ubiquitous. Also, the arguments that convinced us are often based on discussions of cases where someone knows everything there is to know of the form “The so-and-so is F” but does not know which one of the various so-and-sos they themselves are. And it is controversial what exactly we should say about these cases. The net effect has been to make centered worlds and why we need them appear subjects of mystery and controversy to some. Whereas, I want to say, we need them for something as everyday as what happens when we have a perceptual experience, and that what we need – that is, the centered worlds – aren’t at all mysterious. Let’s start by looking at the first point.

To have a perceptual experience is to be in a state that represents things as being a certain way: that the phone is ringing in the next room, that the train I am sitting in has just started to move, that the computer I am looking at has gone to sleep, that there is a spider crawling up my leg. In that sense, a perceptual experience carries putative information, and does so as part of its very nature *qua* experience. We need to learn how a sentence in some given language represents things to be; we don’t need to learn how our senses represent things to be.<sup>7</sup> This means that, for any experience, we can distinguish the ways things might be that are in accord with the experience from those that aren’t. Experiences partition ways things might be. But we cannot think of the partition in question as a partition among possible worlds. That’s because perceptual experiences represent how things are from a point of view without representing how that point of view is other than that it is one where things are thus and so vis-à-vis it. When I have an experience as of a cat in front of me, I am in a state that represents that there is a cat in front of a certain point. My state, however, doesn’t represent where that point is. Equally, it doesn’t represent when it is: 15:00 on 5 May 2012

---

6 See, e. g., John Perry, “The Problem of the Essential Indexical,” *Noûs* 13, no. 1 (March 1979): 3–21; David Lewis, “Attitudes *De Dicto* and *De Se*,” *Philosophical Review* 88, no. 4 (October 1979): 513–543.

7 Though of course (i) we may need to learn what can be inferred *from* our experiences, and (ii) training can help us have experiences with improved representational contents. Radiologists learn how to read X-rays, and part of what happens is that their experiences change as they learn. They come to have, as we might put it, “representationally improved” experiences.

or...? Nor does the experience represent who occupies the point. The only thing it “says” about the point is that there’s a cat in front of it. This means we cannot capture the content of the experience with a set of worlds. Perhaps this is immediately obvious, but let’s spell it out by reviewing the three ways one might seek to capture the content with a set of worlds and why each fails.

First, one might use a set of worlds, with the very same point in each having a cat in front of it. But that would treat my experience as saying which point it is that has a cat in front of it, and that’s wrong. My experience doesn’t say which point has a cat in front of it. Second, one might use a set of worlds, each of which has exactly one point with a cat in front of it, but with the point that has a cat in front of it varying from world to world. The idea would be that by varying the point, one removes the objectionable feature of the previous suggestion: the content no longer says which point has a cat in front of it. But that would treat my experience as saying in part that there isn’t any other point in space with a cat in front of it, and that’s wrong. My experience is silent on whether there is more than one point with a cat in front of it. Finally, one might use a set of worlds, each of which has at least one point with a cat in front of it, and with the point or points that have a cat in front of them varying from world to world. But this would treat my experience as representing that there is at least one cat somewhere or other, and my experience does more than this. It represents that there is a cat in front of a particular point but without representing which point it is. (There is a big difference between my saying that there is a cat in front of me and my saying that there is a cat somewhere or other.)

How then do we get the content of the experience right? By thinking of the experience as partitioning the *points* in a world into those that are the right way in that world to be as the experience represents things to be versus those that aren’t. For each and every point in each and every possible world, it will or won’t be the case that the point has a cat in front of it. Now take every pair of point and world such the point is the right way in the world, then the resulting set is the content of my experience. And this set is, of course, a set of centered worlds.

I have just told you why we need centered worlds to capture the content of one particular experience but I trust it will be obvious that the same argument could be made for any experience. It is a general point about perceptual experiences that we need centered worlds

to capture how they represent things to be. What's more, I hope I have de-mystified centered worlds: they are simply point in a world–world pairs. All the same, there's an issue that needs attention. If experiences have centered content, what makes – what could make, one might ask – an experience veridical? That's the topic of the next section.

### When Is an Experience Veridical?

IF THE CONTENT OF AN EXPERIENCE were a set of worlds, it would be easy to say when that experience is veridical. It is veridical if and only if the set of worlds that is its content contains the actual world. But what should we say given that the content is a set of centered worlds? The question is especially pressing in light of the fact that veridicality is a matter of how the actual world is. Veridicality, like truth, supervenes on the actual world. The actual world determines without remainder what's veridical and what's true, and the actual world is *not* a centered world.

I think that there is only one way to answer this question. It is to appeal to the fact that experiences are had by sentient creatures with a location in space-time. When an experience represents that things are thus and so vis-à-vis some point but without saying where the point is or when it is, the person having the experience, so to speak, locates the point. If you know that I am having an experience as of a cat in front of me and you know where I am and when the experience happens, you know where the cat is likely to be and at what time. Here's an analogy. A photograph represents how things are from a point of view but says nothing about the point of view over and above its being one where things are thus and so vis-à-vis it. But we all know how to ascertain whether or not the photograph is "truthful": find the point from which the photograph was taken and when it was taken – and that is a matter of locating the physical structure that took the photograph and the time the photograph was taken. The camera's taking of the photograph does the locating. If things are the right way in front of that camera when the photograph was taken, the photograph tells the truth. I am saying the same about an experience and the person who has the experience.

So, in sum, when is an experience  $E$  veridical? Just when the set of centered worlds, the set, that is, of point in a world – world pairs, that is  $E$ 's content, contains a pair consisting of the point in space where the creature having  $E$  is located at the time of having the experi-

ence and the actual world. So, to say it in English (more or less), my experience as of a cat in front of me is veridical just if, in the actual world, there's a cat in front of where I am located at the time I have the experience. (The fact that this sounds platitudinous is of course confirmation.) Here I am fudging something. It is a nice question exactly how the creature that has the experience locates the point from which its experiences represent things to be (though this is not, as we've highlighted, part of the content). Is it the creature's body that does the locating, or is it some part of the brain's location that's crucial? Moreover, surely we shouldn't be presuming materialism in this discussion. We want an account of how the creature that has the experience locates the relevant point of view which makes good sense in a Cartesian world.


I don't want to minimize the importance of these questions but I do think they are fairly classified as questions for another time. Any plausible answer to them can be plugged into the account of when an experience is veridical that I have just given.<sup>8</sup>

## From Experiences to Sentences

AN AWFUL LOT OF THE SENTENCES we produce by way of making claims about how things are come from our perceptual experiences. Sometimes the path is a pretty direct one. I feel hot and say in consequence "I am hot," or, maybe, "The room is hot." Sometimes the causal chain is more complex and involves a number of (hopefully) information preserving links. Someone at a conference hears an interesting paper on climate change. They send an e-mail to a friend. The friend passes the news onto Jones and she acquires a view about climate change as a result. At each stage, a combination of the content of various experiences, views about the contents of various sentences that are represented as appearing on computer screens or as heard etc. plus collateral information (which will in its turn typically come from experiences in one way or another) culminates in Jones producing sentences that represent how things are in regard to climate. The point is too obvious to need further elaboration.

---

8 I am also ducking an important, related issue. One can have a veridical experience as of a single cat's being in front of one – there really is one and only one cat in front of one – without its being the case that one sees that cat. Examples that go back at least to H. P. Grice, "The Causal Theory of Perception," *Proceedings of the Aristotelian Society: Supplementary Volume* 35 (1961): 121–168, tell us this.




But it tells us something that is perhaps not so obvious. It tells us to expect that the contents of very many sentences will be given by sets of centered worlds, and not by sets of worlds. Indeed, we should expect sentences that have contents that can be given by a set of worlds to be the exception.

I know this worries some people. It can appear to make sentences in general too self-regarding. When I produce a sentence with centered content, aren't I, in one way or another, saying how things are with me? I think that there is a sense in which this is true, but it is not a sense that should worry us. It is exactly what we should expect. Take a sentence like "There is a cat in front of me." It is plausible that the content of that sentence is given by a set of centered worlds, each center of which has a cat in front of it (at the time in question). As we have seen, that kind of account is right for the corresponding experience, and it is hard to believe that things would be different for the sentence that reports the experience. But it does not mean that the content of "There is a cat in front of me" is a set of centered worlds with a cat in front of the person producing the sentence. The sentence says nothing about producers of sentences or the sentences they produce. It isn't about a sentence or the producer of a sentence; it is about a cat. The role of the producer of the sentence is to locate the center for those of us who come across the sentence. The producer of the sentence is a center locator, as I will call it.<sup>9</sup> If you want to know where the cat is, attend to where the producer of the sentence is and work from there. Knowing this is part of knowing how "me" works in sentences of this kind. But it isn't part of how the sentence represents things to be; that is why the producer of the sentence isn't part of the content. Perhaps it is worth laboring this point by spelling out how information gets acquired when we come across sentences with centered content. Here is how it goes for our example.

1. You come across a token of "There is a cat in front of me."
2. You have reason to trust it (you have reason to think that I'm not lying, my eyes are working properly etc.).
3. Its content is a set of centered worlds with a cat in front of their centers.

---

<sup>9</sup> And see Jackson, *Language, Names, and Information*. The role of the producer of the sentence in locating the center for the sentence parallels the role of the person's having the experience in locating the center for the experience.



4. In understanding the sentence, you know that this is its content.
5. In understanding the sentence (more especially, the way “me” works in this context), you know in principle that the center is the producer of the sentence.
6. You know that you and the token sentence are in the same world.
7. You know, or can know, how you are related to the producer of the token sentence and thereby your relationship to the center in question.
8. You learn where the cat is to be found, as you know it is in front of the center in question.

At no stage in the above derivation do we suppose that the sentence or the producer of the sentence is part of the content. Of course, the sentence and its producer play a role in the account of how we extract information from coming across the sentence. But that is as things should be and so is no objection. We all agree that the information we get from a telescope is in part a function of its location and orientation; the same goes for utterances and inscriptions.

We mentioned earlier the need to detail what it is for an experience with centered content to be veridical. Correspondingly, we need to detail what it is for a sentence with centered content to be true.<sup>10</sup> In the case of a sentence whose content is a set of worlds, we can say that it is true just if the set of worlds that is its content contains the actual world. But we cannot say this for a sentence whose content is a set of centered worlds. We need, in the case of these sentences, to bring the sentence itself, the token, into the mix, much as we brought the experience itself into the mix when discussing veridicality. Each sentence with centered content locates its center. In the simplest case, its center is the producer of the sentence. Now we can say when a sentence with centered content, *S*, is true: it is true just when the set of centered worlds, the set, that is, of point in a world – world pairs, that is *S*'s content, contains a pair consisting of the point located by *S* and the actual world. So when is “There is a cat in front of me” true? When there is a cat in front of the producer of the sentence in the actual world. To echo an earlier remark, this sounds almost platitudinous, and I think that's confirmation.

---

<sup>10</sup> The importance of doing this is highlighted by Robert Stalnaker, *Our Knowledge of the Internal World* (Oxford: Oxford University Press, 2008), 50. I am not suggesting that he would agree with what I say.

## Back to the Puzzle About “Water”

NOW WE ARE ABLE TO SOLVE the puzzle about “water.” How come the word “water” in the early 1700s was used to give information about the watery kind and yet it fails to refer XYZ on Twin Earth; how come XYZ isn’t the right way for “water” to apply to it, despite its being every bit as watery as H<sub>2</sub>O?


I think we use the word “water” for the watery kind *we have often come across*.<sup>11</sup> The kind isn’t always liquid and all that; it isn’t always, as we might put it, especially watery. But it is the kind which is very often liquid and all that, and which we have encountered. If that is how the word was used in the early 1700s, it is transparent how come it was a word for giving information about the watery kind. It is also transparent why the word (in our mouths) doesn’t refer to XYZ on Twin Earth. XYZ isn’t the kind we Earthians encounter; it’s the kind the Twin Earthians encounter. On this view, sentences containing “water” have centered content. The word is a word for talking about a kind we have come across, a kind that stands thus and so vis-à-vis ourselves. The content of, for example, “Water is plentiful” is, accordingly, the set of centered worlds with centers that have encountered a watery kind that is plentiful.

## The Interplay between the Set of Centered Worlds and the Center Location Rule

I HAVE SAID that there are lots of sentences with centered content. I have given a few examples but I trust it will be obvious how to generate many more. As I said earlier, having centered content is the norm. I also said that, in order to get information from sentences with centered content, you need to know how sentence tokens locate the centers in question. I gave one example: the way the appearance of “me” in certain contexts tells one that the center is the producer of the sentence. An obvious topic for future work is giving accounts, for a whole range of sentence constructions that have centered content, of how tokens of these sentence constructions serve to locate the relevant centers. This lecture is not the place for

---

11 The “we” are those of us who use “water” in such a way that Putnam is right that XYZ on Twin Earth doesn’t fall within its extension. Of course a version of this idea is to be found in Putnam, “The Meaning of ‘Meaning.’”



this exercise. However, it is, I think, worth noting how an account of the set of centered worlds for a sentence can intersect with an account of the how the token sentence locates the center.

Take the sentence “It is hot 1km due North of me.” One way to give its content is as the set of centered worlds with centers 1km South of a hot region. The corresponding center location rule will then be: the producer of the sentence token. However, one could equally do it this way. The content is the set of centered worlds with hot centers, and the location rule is: 1km due North of the producer of the sentence token. Surely, it is obvious that there is no real difference between the two ways of spelling things out. I am sure that this is not an isolated example. There will many cases where there are two different but equivalent ways of skinning our cat. Shortly, I will offer a short account of the informational value of proper names in terms of a set of centered worlds and a center location rule. I am sure that there are different ways of specifying the set of centered worlds combined with correspondingly different location rules that deliver the same net outcome for the informational value of proper names.<sup>12</sup>

## Proper Names


PEOPLE WHO UTTER SENTENCES of the form “*N* is *F*” give information about how things are, often information properly speaking, as is the case for those who utter “Paris is pretty” or “Plato was a great philosopher.” You might think of the description theory of reference for (proper) names and the direct reference theory of names as offering competing accounts of this information (there are other ways of thinking of those two theories of course). The description theory might be thought of as saying that the information is given by a set of worlds where the so-and-so is *F*, and the challenge set by Saul Kripke’s attack on the description theory is then to come up with a value for “so-and-so” that survives his objections.<sup>13</sup>

---

12 Having said this, recent conversations with Dan Marshall have impressed on me the advantages of adopting as an invariable center location rule: the producer of the sentence token, and making the appropriate adjustments to the set of centered worlds. I have, however, not followed this policy in the section that follows.

13 Saul Kripke, *Naming and Necessity* (Cambridge: Harvard University Press, 1972).






The direct reference theory might be thought of as saying that we should give up on trying to save the description theory and instead affirm that the information is given by the set of worlds,  $w$ , where  $N$ , that very thing and independently of what descriptions it satisfies in  $w$ , is  $F$ . From the perspective of this lecture, both suggestions are highly dubious. The ubiquitous role of experience in grounding our claims about how things are, including those we make using names, strongly suggests that we should be looking for an account in terms of a partition in centered worlds, not in worlds – be they specified à la the description theory or à la the direct reference theory – combined with an account of how names serve to locate the relevant centers.

Here is one way to do this. Take the set of centered worlds for “ $N$  is  $F$ ” to be the set of centered worlds whose centers are  $F$ . Here we take a leaf out the direct reference theorists’ book in the sense that there is no descriptive requirement being placed on the centers by “ $N$ .” Then think of the token name as serving to locate those centers. How might it do this? Think of how historians approach questions like “Did Helen of Troy really exist?” and “Are the Robin Hood stories about a real person, and if they are, was he the good guy the stories make him out to be?” Historians approach these questions by looking for information-bearing causal chains that run from tokens of the names in various texts back in time to persons named in such a way that they generated those chains. In doing this, they treat names as building blocks in information highways. And this is indeed what we do every day with names. This is how names get to be wonderful devices for passing on information about individual objects. That indeed is the underlying rationale for giving objects names in the first place.<sup>14</sup> Here we are taking a leaf out of the causal theory of reference for names. What we have just been saying is not that different from the kinds of things one finds in *Naming and Necessity*, and in the writings of some of those influenced by this book. (I am not saying that this is how they think of matters.)


There is, however, one important difference. Many supporters of the causal theory of reference for names insist that the causal facts that secure the reference of a proper name are not common knowledge; that’s typically how they distinguish themselves from supporters of

---

14 For more on this (but the basic idea can be found in many places, as you would expect given that it is folk theory), see Jackson, *Language, Names, and Information*, ch. 5.



causal descriptivism. But this sits ill with the way the method is used by historians to answer questions about Helen of Troy and Robin Hood. They don't feel the need to get advice from philosophers of language. Or consider reviewers of their books. These reviewers express opinions about whether or not the authors have or have not shown that Robin Hood or Helen of Troy or Job or . . . really existed and did this, that or the other thing. But mostly the reviewers are not experts in the philosophy of language; they are historians and writers, and moreover they take it for granted (rightly) that both the books they are reviewing and their reviews of them are accessible to reasonably intelligent readers quite independently of whether or not they have done philosophy at university. And this is to take it for granted that reasonably intelligent readers know what it takes to be the right kind of causal origin of the tokens of "Robin Hood" and "Helen of Troy," or, if it comes to that, "Tony Blair." This is what I had in mind when I said earlier, in a footnote, that I was offering a bit of folk theory.



## Acknowledgments

I was delighted to be asked to give this lecture. In writing the lecture up for publication, I have taken account of the helpful discussion at the lecture. I have also drawn on many discussions over many years, most recently after oral presentations of related material at Cambridge University, May 2011, and Adelaide University, March 2012.

## References

- Grice, H. P. "The Causal Theory of Perception." *Proceedings of the Aristotelian Society: Supplementary Volume* 35 (1961): 121–168.
- Jackson, Frank. *Language, Names, and Information*. Chichester: Wiley-Blackwell, 2010.
- Kripke, Saul. *Naming and Necessity*. Cambridge: Harvard University Press, 1972.
- Lewis, David. "Attitudes *De Dicto* and *De Se*." *Philosophical Review* 88, no. 4 (October 1979): 513–543.
- Perry, John. "The Problem of the Essential Indexical." *Noûs* 13, no. 1 (March 1979): 3–21.
- Putnam, Hilary. "The Meaning of 'Meaning.'" In *Mind, Language and Reality*. Cambridge: Cambridge University Press, 1975.
- Stalnaker, Robert. *Our Knowledge of the Internal World*. Oxford: Oxford University Press, 2008.