

Self-notions

John Perry

(Published in *Logos*, 1990: 17-31.)

December 13, 1998

[In *Logos*, 1990, 17-31.]

Abstract

"Self-beliefs" are beliefs of the sort one ordinarily has about oneself, and expresses with the first person. These contrast with the beliefs one has in "Castañeda cases," in which one has a belief about oneself without knowing it. This paper advances an account of the nature of self-belief. According to this account, self-belief is a special case of interacting with things via notions that serve as repositories for information about objects with certain important relations to the knower, and as motivators for actions the success of which is dependent on the object in that relation to the agent. Identity is such a relation, and "self-notions" play this special role: they are the repositories for information gained in normally self-informative ways, and the motivators of types of action whose success normally depends on facts about the agent. Self-beliefs involve such self-notions, while the beliefs that one has about oneself in Castañeda cases do not.

"Not long ago," Mach wrote in 1885, "after a trying railway journey by night, when I was very tired, I got into an omnibus, just as another man appeared at the other end. 'What a shabby pedagogue that is, that has just entered,' thought I. It was myself: opposite me hung a large mirror. The physiognomy of my class, accordingly, was better known to me than my own."¹

Mach acquired a belief at the beginning of the episode, which we can imagine him expressing as:

(1) That man is a shabby pedagogue.

By the end of the episode he has another, which we can imagine him expressing as:

(2) I am a shabby pedagogue

¹Ernst Mach, *The analysis of sensations*, translated by C.M. Williams and Sydney Waterlow, (Chicago & London: Open Court, 1914), p. 4n.

While (2) expresses self-knowledge, (1) does not. And yet (1) is, in a perfectly clear sense, a case of Mach believing something about himself. Mach implies that he was, in fact, a shabby pedagogue. It is because he was a shabby pedagogue that we take the belief expressed by (1) to be true. If it is Mach's being a shabby pedagogue that makes (1) true, then (1) was about Mach, and expressed a belief about him. Nevertheless (1), unlike (2), is not an expression of self-knowledge.

Henceforth I shall use the unfamiliar term "self-belief" rather than "self-knowledge." Although "self-knowledge" sounds more familiar, it is somewhat misleading since the distinction between knowledge and mere belief is orthogonal to the issues I discuss. I take beliefs to be complex cognitive particulars, that come into existence as a result of perceptions, inferences and other events, and influence the occurrence and nature other beliefs and actions. I assume that two beliefs are involved here, one that Mach acquired when he stepped on the bus, and one that he acquired a bit later, when he figured out that he was looking at himself. I want to understand the difference between those beliefs. It is not sufficient, for this purpose, to note that (2) contains the word "I" where (1) contains the words "that man". This is why (2) is an expression of self-belief and (1) is not. But I want to know why the belief thus expressed is a self-belief.

1 I

My account will presuppose a fairly common sense view of beliefs. The mind has ideas of things, properties and relations. I'll call ideas of things notions. A belief is a complex in which an idea of a property or relation is associated with the appropriate number of notions of things. The content of a belief is that the things the notions represent stand in the relation that the idea represents. A number of beliefs with a common notion is a file.

The function of beliefs is to retain information picked up through perception, to formulate hypotheses, to allow for the combination and comparison of beliefs and the formation of new beliefs through inference and to motivate actions that will promote one's desires if the beliefs are correct. The objects an idea represents depends on its role in psychological system within which beliefs play their role and the way that system is set into the wider world. When Mach stepped on the bus, he formed a notion of the man he saw stepping on the bus. The fact that this was a notion –an idea of a thing – and not some other kind of idea depends on the way it functions internally; the fact that it was of a certain man (Mach, as it turned out) depends on external circumstances. If Mach had been looking at some other man, the very same idea would have been a notion of that other man.

We have unlinked notions of the same thing when the external factors that determine which thing the notions represent happen to make them represent the same thing, although there is nothing in the notions themselves, or the ideas associated with them, that reflects this identity. The beliefs involving those different notions of the same thing can function independently. They can arise at different times, and can affect actions in quite different ways. All of the "morning star/evening star" cases from the philosophy

of language are instances of this.²

Although it is theoretically possible for there to be two unlinked notions of the same object, which are associated with the very same ideas, this seems unlikely. To have two unlinked notions of the same thing, we must have interacted with it in different ways or circumstances, or at least at different times, and some of the associated ideas are likely to reflect these differences. Suppose I have two unlinked notions of you, one formed as a result of reading articles by you, one formed as a result of seeing you in the library of my university, where you happen to be visiting without my knowing it. When I read your latest article, the first notion will become associated with new ideas based on what I read. When I see you in the Student Union, the second notion will become associated with new ideas based on what I observe. I have two clusters of beliefs, or files, about you, each consisting of all the ideas associated with one of the two notions.

Beliefs and files about the same object may motivate quite different actions. When I interact with you, my behavior will be guided by one file or the other, depending on the situation. I know how to write to you, as the author of the articles I have read, for they include your name and a department where you can be reached. Perhaps I write you a letter praising you for your sensitivity and clarity of thought. I know how to speak to you, as the person I have seen in the library, for the file corresponding to my first notion contains information about where you hang out and what you look like. Perhaps, based on my observations of you around campus, where I have seen you go out of your way to be kind, I also think you are sensitive. This leads me to associate the first notion with the idea of sensitivity. This doesn't lead me to write you a letter. There is no information in this file about your name or address. I just walk up to you one day and compliment you on your kindness and sensitivity. As long as the files remain unlinked, the information in one will not affect the actions guided by the other. For example, I won't call you by your name when I see you.

In this example, I have two beliefs with the same content: that you are kind. The beliefs have quite different causal roles. This is explained by the different notions involved in the beliefs, and the different ideas that those notions are associated with. While the beliefs have the same content, the files of which they are a part do not. What unifies a file, and makes two beliefs about the same person relevant to one another, is not that they are about the same person, but that they contain the same notion or linked notions. One could have a file made up of a lot of information about different individuals mistakenly associated with a single notion, just as one can have two notions where there is only one object.

The examples of this sort of thing that we find in the philosophy of language often involve rather curious episodes from the history of thought, such as the Babylonians having two notions of Venus, or unusually confused individuals, such as Kripke's Pierre, – or Mach on the bus, for that matter. But the phenomenon itself is very common and doesn't require unusual circumstances or unusually confused people. I see

²See Jon Barwise and John Perry, *Situations and Attitudes*, especially pp. 248ff; Mark Crimmins and John Perry, "The Prince and the Phone Booth: Reporting Puzzling Beliefs," *The Journal of Philosophy*, LXXXVI (1989): 685-711; Mark Crimmins, *Talk About Beliefs*, Ph.D. dissertation, Department of Philosophy, Stanford, 1989.

my friend Al limping towards me, but cannot yet recognize him; I form a notion of this person. At that moment I have two unlinked notions of Al. Certain of my beliefs about Al I have twice over, such as that he was a man. Others I had in one file but not in the other, such as that he had a limp. I accumulate information about him as he gets nearer; finally I recognize him as Al. At that point the notions become linked; the newly acquired perceptual information combines with the old information and I say, "Why are you limping, Al?" If the identification is tentative, the notions may retain their identity; if not, they may merge and become one.

2 II

Mach's confusion is a special case of this sort, where the person he comes to recognize is himself. At the beginning of the episode, Mach had two notions of Mach. One he acquired when he stepped on the bus and saw what he took to be a man at the other end. The other is a self-notion, the sort of notion that is usually involved in his beliefs about himself. Beliefs involving one's self-notion have a special role in one's cognitive life, and we usually reserve the term "self-knowledge" for knowledge involving beliefs of this sort. But we need to say more about this special role.

The natural place to look is the ideas with which self-notions are associated. Consider the self-idea, – the idea we would express as "being me". The notion involved in Mach's first belief isn't associated with this idea, while the one involved in his latter belief is. But which idea is this? We cannot identify the idea by the property it represents. Mach has two ideas of the property of being Mach, one which he would express with "being me" and one which he would express, directing his attention towards the man he sees, with "being him". The former is the self-idea, but why? What makes one of the ideas that represents the property of being identical with Mach his self-idea, and not the other? We might suppose that the self-idea is a complex idea, composed of the idea of identity in association with the self-notion. This seems plausible, but now we have just gone around in a circle: what is special about the self-notion is that it is associated with the self-idea; what is special about the self-idea is that it has the self-notion as a constituent.

Another approach is to characterize self-notions semantically. We can think of a special self-notion on the analogy of the indexical "I". Just as utterances of "I" stand for their utterers, these special self-notions would stand for the thinkers to whom they belong. Such notions, and any other notions that were linked to them, would be self-notions. Beliefs with self-notions as constituents would be self-beliefs. The notion Mach acquires when he looks at the fellow in the mirror is not a self-notion at the beginning of the episode, but becomes one when he recognizes himself, and links his "that man" notion to his indexical self-notion. This characterization leaves a question unanswered, however. Consider Mach's two notions of himself. Which of these should be characterized as a self-notion? They both represent Mach, so that can't be the difference between them.

An analogous question about language would be why we take "Ernst Mach" to

be a name of Mach, and “ich” to be the first person in Mach’s idiolect. They both stand for Mach, so that isn’t the difference between them. The answer to this question lies in the different ways Mach uses these terms in communication, both as speaker and as listener. He takes every utterance containing “Ernst Mach” to be about him, while he takes utterances containing “ich” to be about the speaker.³ This is suggestive, but doesn’t quite get at the difference between Mach’s two notions of himself. These notions are not devices for communication. They are not public; Mach does not produce them as an aid to securing recognition by others of his communicative intentions. Rather, they are parts of a system for the pick-up, retention, analysis, discovery and utilization of information by an individual. To explain what we mean by self-notions, and how Mach’s two beliefs differ, we need to explain the role of these notions and beliefs in this informational system.

3 III

Just as there is a special way of thinking about the person you are, there is a special way of thinking about the place you happen to be in, the way you think of the place you call “here”.⁴ Without realizing it, I could be in, say, Grand Island, Nebraska, at the same time I was watching a video of Grand Island scenes. When I thought, “that city looks like a fun place to visit,” as I watched the screen, I would be thinking about Grand Island, but not as the place I was at, as “here.” Thinking for a bit about what is involved in thinking about a place as “here,” will be helpful in seeing what is involved in thinking about a person as “I”.

Suppose you are traveling, and wake up in a hotel room in Grand Island. You look out the window and see rain. So you grab your umbrella before departing from the hotel. Here is a simple case of using information acquired perceptually to guide action. When you look out the window, you get information about the weather in Grand Island. And when you depart from the hotel, your decision to take an umbrella is vindicated because of the weather in Grand Island. The fact that the place whose weather you learn about when you looked out the window is the place whose weather determines whether you need an umbrella is crucial to the success of your use of the information. How do you have to think about Grand Island, to facilitate this use of information to guide action?

One possibility is that you think about Grand Island via its relation to you. You think “It is raining here” or “It is raining in this city”. A second is that you think about Grand Island via some attribute that is independent of its relation to you, such as its name. The appropriate expression of your thought is “It is raining in Grand Island.” These different ways of thinking correspond to different beliefs you might hold, independently of one another. A person who doesn’t remember which city he is in, looks out the window and sees rain, will have the first without the second. If instead

³I am ignoring the complication provided by other people named “Ernst Mach”.

⁴I am ignoring the complication that “here” isn’t usually a name or pronoun referring to a place but an adverb of place.

he turns on the news and hears the reporter on the Lincoln station say “It is cloudy in Omaha, sunny in Lincoln, and raining in Grand Island,” he will have the second without the first.

Let’s imagine that you have a very poor memory, and keep track of information by using 3x5 cards. You have a number of these for the various cities you frequent, including one for Grand Island. When you hear the news report you take out this card and jot “rain” on it. But suppose you have forgotten where you are. You don’t know which of the city-cards to write it on. In this case, you might take out a new card and jot “the place I’m in; rain” on it. Call this card the “here-buffer”. You might keep a card for this purpose, with a permanent entry, “the place I’m in”. It would be a cousin of the indexical “here” in “3x5 card-assisted thought”. As information accumulated on this card, you may be able to figure out where you are. When you do, you might transfer the information from the here-buffer to the Grand Island card. At that point the cards would not only in fact contain information about the same place, but be recognized by you as doing so. They would be linked.

As time passes, you have to update your cards in various ways. But the relatively permanent features you have noted on your Grand Island card, such as “on Interstate 80”, “has an interesting museum” do not need to be changed just because you move on to Kearney or North Platte. This card is of Grand Island, whether you happen to be there or not. But your “here-buffer” should be erased, not because facts have changed about Grand Island, but because it has ceased to be the city you are in. And – assuming you are aware of the fact you have moved on – it will be unlinked from your Grand Island card.

Let’s call a notion a “here-notion” if it is associated with a self-notion and the idea of being-in. For the person who doesn’t know where they are, the only here-notion will be a here-buffer. Note that this way of having the information about rain in Grand Island suffices to motivate taking an umbrella. What one needs to know, to determine whether rain in a given city provides a reason for taking an umbrella is whether that city is the city one is in. You know that the city you see out your window is the one you are in, without knowing which city it is. On the other hand, having the information that it is raining in Grand Island via a notion that isn’t a here-notion and doesn’t reflect that fact that it is the city you are in, will not suffice to motivate taking an umbrella, however comprehensively the beliefs with that notion as a constituent may cover other issues about Grand Island.

The reason for this is that there are ways of getting information about the city one is in quite independently of which city it is. And there actions the success of which depends on the conditions prevailing in the city one is in, quite independently of which one it is. Looking outside one’s window is a way of finding out what the weather is like in the city one is in, whichever city that happens to be. And taking an umbrella is an action that will be a good idea, if it is raining in the city one is in, whatever city it happens to be. I shall say that there are “normally here-informative” ways of getting information, and “normally here-dependent, here-directed and here-effecting” ways of acting. It will be reasonable for normally-here informative ways of getting information about cities to motivate normally here-dependent actions whose success depend on that

information. This is what happens in the case in question.

I shall call relations between an agent and another object – including places, material objects and other persons – that support such special ways of knowing and acting, an “epistemic/pragmatic” relations. The relation of being in, a relation between people and places, is an epistemic/pragmatic relation. There are many others. There are special ways to know about the material objects and people in front of one (open your eyes and look, reach out and touch), and special ways of dealing with them. There are special ways of knowing about the person who wrote the article you are reading, and special ways of communicating with them. There are special ways to know what a person is saying when they have called you on the phone (listen to the sounds coming out of the ear piece) and special ways of saying things to them (speak into the mouthpiece). Where R is an epistemic/pragmatic relation, we may speak of “normally R-informative ways of perceiving” and “normally R-directed/dependent/effecting ways of acting”.

We are all masters of hundreds of such ways of gaining information about things, and dealing with things. They allow us to gain information about things and deal with things without having any way of identifying them independently of their relation to us. They allow us to interact with individuals we know about, once we determine or bring it about that they stand in an epistemic/pragmatic relation to us. When you call, I use such methods to accumulate information about who is talking to me, until I figure out that it is you. Once I realize it is you, and have combined the information in my “on the phone” buffer with that associated with my “objective” notion of you, I can use all of this information to decide what to say to you, which I do in the way that one says things to the person with whom one is talking on the phone.

The informational role of an R-notion is to serve as the normal repository for information gained in normally R-informative ways, and as the normal motivator for normally R-effecting and R-dependent actions. The information I pick up by looking around me will, normally, become associated with my here-notions. The beliefs involving these notions will motivate actions like taking an umbrella, whose success depends on the weather around me. The information that motivates a normally here-dependent action need not have been obtained in a normally here-informative way. You may take an umbrella because you heard on the radio that it was raining in Grand Island. The action motivated by information gained in a normally here-dependent way may motivate actions whose success is not normally here-dependent. Seeing that it is raining, you may tell someone on the phone that it is raining in Grand Island, a statement that would be as true if you said it standing in Valentine or Ainsworth.

There are also ways of acting that involve having information via objective (or at least more or less objective) notions. By an objective notion, I mean one that is associated with identifying properties of an object that do not involve its relation to the agent who has the notion. An R-notion can also be an objective notion. If one knows that one is in Grand Island, for example, one will have a single notion that is associated with such ideas as “being the Nebraska city called “Grand Island”” as well as with “being the city I am in”. When we want to combine information gathered in different circumstances about the same thing, we need to get at the information in ways that are not tied to one or the other circumstance. If we want to use this information from the

“permanent file”, we need ways of recognizing the object the file is about. Consider what happened, in the example above, that enabled me to ask “Why are you limping, Al?” I had to recognize the person limping towards me as Al; that is, I needed to think of him not only via the relation of being in front of me, but objectively, as the person who had a certain appearance and manner (whether I am looking at him or not) and who is called “Al”.

“Not knowing who someone is” usually amounts to having an R-buffer and a permanent file that are unlinked. There are two sorts of cases. In the example above about my friend Al, I had the buffer and need to pick the right file. But suppose instead that Al is a philosopher I have read and written to; he is in a room full of philosophers and I want to talk to him, but don’t know what he looks like. I’ve got the file, but need to pick the right buffer. A calendar entry, with a date and some appointments by it, is like a permanent file. It doesn’t help if you don’t know what day it is—if you don’t know whether to think of the day as “today” or “tomorrow” or what. But it would be equally frustrating to be in the position the cartoon character Ziggy was once depicted as being in: he rips off one page on his calendar and reads “the next day” on the next page. We want calendars to give up objective representations of days so that we can use them to store and combine other information given objectively.

We might think of our notions as forming a multi-leveled system. At the top level are notions that are completely, or at least maximally, independent of relationships to us. These are “objective” representations. The lower levels contain buffers for various relationships to us, associated with various epistemic/pragmatic relations, of increasing specificity. We pass information up the levels, as we gather information about objects in epistemic/pragmatic relations to us, recognize them and store the information in ways that are more independent of our relationships. We pass information down the levels when we recognize an object and act on it ways that depend not just on its present relation to us, but other properties about which we have gathered information in the past.

When we think of beliefs, we are usually thinking about information stored at the relatively high levels. In fact, it is difficult to describe links between levels if we confine ourselves to “believes” and its cognates, as any philosopher who has dealt with the puzzles from the philosophy of language is aware. We have an additional vocabulary, including “recognizes,” “takes to be,” and “identifies” to describe linking. For objects and persons with which we are familiar, we have relatively rich permanent files, and it is the contents of these files that primarily count as our beliefs about the thing or person in question. Such beliefs provide the extra or incremental information we have to bring to bear on our interactions with these objects and persons, in addition to what we perceive about them at the time of a given interaction.

4 IV

I believe what is special about self-notions is that they are the normal repository of normally self-informative ways of perceiving, and the normal motivator of normally

self-dependent ways of acting.

We might call the example about Mach with which I began a "Castañeda example, after Hector-Neri Castañeda, who introduced a number of examples of this sort, and insightfully analyzed them."⁵ They typically involve perceptual states that are not normally self-informative in the sense I am using the term. The state Mach was in, when he saw the man in the mirror, was the sort of state one is usually in when one sees that someone else, standing at some distance from one, is shabbily dressed, not when one sees oneself to be shabbily dressed.

Contrast with such cases, what we might call Shoemaker cases. Sydney Shoemaker has emphasized that we often find things out about ourselves in ways that are "immune to mis-identification".⁶ Suppose you are at a party. You bend over to pick something up, and hear the ripping sound that is characteristic of trousers splitting. Then you feel a hot flush in your face. You are aware that you are blushing. Now who is it, of whose blushing, you are aware? We are almost inclined to say that the question makes no sense. It is of course your own blushing of which you are aware. It's not that you cannot be aware of the blushing of others. You can see them blush. But you can't feel them blush; you can't come to know someone else is blushing, in the way that you typically come to know that you are.

Shoemaker emphasizes that immunity to mis-identification should not be confused with incorrigibility or even privileged access, although they often go together. Compare blushing with being embarrassed. It seems that there is a way of knowing that one is embarrassed, the normal way, which is immune from mis-identification, privileged, and at least close to incorrigible. In the pants-splitting episode, I can't be wrong that I am the one that is embarrassed (immunity); I know this in a way that is more direct and error-free than anyone else can (privilege), and perhaps I can't be wrong about it (incorrigibility). But I can be wrong about whether I am blushing. I may know that I am embarrassed, but be mistaken in thinking that I am blushing. I may not be in as good a position to tell if I am blushing as someone else, who can see my face redden. So my judgement that I am blushing is neither incorrigible nor privileged. But this judgement cannot be wrong because the person I take to be blushing, is blushing, but is not me. Feeling one's face flush is a corrigible way of finding out that one is blushing; it cannot be used to find out whether someone else is.

It is the way of finding out, not what is found out, that is immune to mis-identification. One could look in a mirror, and think that one saw oneself blushing, although it was someone else. I could believe that I am blushing, and believe it on the basis of observing someone blush, but be wrong about who it is. on it.

These ways of knowing that are immune from mis-identification are, I suggest, just

⁵See especially Hector-Neri Castañeda, "'He': A Study in the Logic of Self-Consciousness," *Ratio*, 8 (1966): 130-57; "Indicators and Quasi-indicators," *American Philosophical Quarterly*, 4 (1967): 85-100; "On the Logic of Attributions of Self-Knowledge to Others," *Journal of Philosophy*, 65 (1968): 439-56. For a discussion of Castañeda's views, see my "Castañeda on He and I," in James E. Tomberlin, ed., *Agent, Language and World* (Indianapolis: Hackett, 1983), pp. 15-41.

⁶See Sydney Shoemaker, *Identity, Cause and Mind* (Cambridge, England: Cambridge University Press, 1984), especially "Self-reference and self-awareness." See also Gareth Evans, *Varieties of Reference*, (Oxford: Oxford University Press, 1982), especially sections 6.6 and 7.2.

a special case of "normally R-informative ways of perceiving". A perceptual state S is a normally self-informative way of knowing that f(y), when the fact that a person is in state S normally carries the information that that person is f, and normally does not carry the information that any other person is. Identity is an epistemic/pragmatic relation. Feeling one's face flush is a way of registering the information that the person identical with the feeler is blushing. Feeling hunger is normally a way of detecting that one's own stomach could use some filling. The feeling of needing to urinate is normally a way of knowing that one's own bladder is full. In each case, someone else can determine the same thing, using a different technique. Perhaps you can see me blush even when I am not aware that I am. Perhaps you know that my stomach is full, having noticed what I have put into it, while I am still in that charming interval between being full and feeling full. Parents often are better judges of how full children's bladders are than the children themselves are. But you cannot (normally) know that I am blushing, or full, or need to urinate in the way that I do.

Why do I say normally? There are some cases where this qualification is clearly required. Think about watching your hands as you type or play the piano. There is a characteristic way of seeing one's own hands and limbs and torso, a way in which one hardly ever sees anyone else's hands or limbs or torso. Yet when we learn to play the piano, or type, our teacher's hands could conceivably be mistaken for our own; one sees them in the same way as one sees one's own hands. Of course one would quickly spot the mistake, since however similar one's teacher's hands are to one's own, one cannot move them like one moves one's own.

Suppose that a way is developed to repair spinal column injuries by using an external shunt, which connects the column below the injury to the brain stem. One can imagine the shunt having an external connection for some reason. One could go on to imagine that there was enough similarity among people that one person's shunt could be plugged into another person's brain stem receptor so that coherent signals would arrive at the latter's brain about the bodily conditions of the former. (In fact, although this seems a possibility, I don't think we have any reason to suppose it is a very likely one.) When a fly landed on the first person's leg, the second person would feel the sensation, and perhaps slap his own leg. The second person would be perceiving that there was a fly on someone's leg, in the way that normally perceives that there is a fly on one's own leg. Because of possibilities like this, I add the qualification "normally".

Paralleling normally self-informative ways of knowing are normally self-effecting ways of acting. Towards the end of the movie *Spellbound*, we see Leo G. Carroll point his gun at Ingrid Bergman as she walks out the door. Holding the gun in this way and pulling the trigger is a way, in the circumstance in which there is a person in front of the agent, of killing that person. As the movie continues, we Bergman continue to walk away, toward the door of Carroll's office, from his perspective. Slowly, we see the hand holding the gun turn, until the barrel of the gun is all that is visible on the screen. Then it fires. We know what Carroll has done, and to whom. He has killed someone, and the someone is him. The way Carroll held and fired the gun was a normally self-effecting way of killing someone. Of course, if Carroll had a head shaped like a donut, he could have shot the someone behind him. But normal people normally kill themselves when

they shoot like that.

This is only a particularly dramatic case of a whole class of actions. Imagine George and Barbara Bush seated across from each other at a boring dinner. Both know that the president is thirsty. Both may desire that he get a drink. The appropriate action for the President to take is the familiar one of reaching out and bringing the glass of water towards his lips. That is an action that will succeed if the agent is thirsty. It is a normally self-dependent/directed/effecting action. It won't do any good for the First Lady to perform it. At least, it won't help relieve the President's thirst. She should pick up the water and offer it to him. That is a way of relieving (or helping to relieve) the thirst of someone sitting across from you.

I suggest, then, that self-notions are those that have the special causal role of being the repository for normally self-informative ways of perceiving and the motivator of normally self-dependent actions. Just as one can see that it is raining where one is and decide to take an umbrella, without knowing which city one is in, one can gain information about oneself and apply it without knowing who one is. A thirsty amnesiac, who doesn't know who he is, or a character from a Dickens novel, who may have never known, can still drink glasses of water, eat when hungry, and the like. Normally self-informative perceptions can trigger normally self-dependent actions without needing to be linked to any self-independent notion of oneself.

There is one big difference between identity and most other epistemic/pragmatic relations. With most of the others, a given agent will stand in the relation to different objects at different times. The place one is at, the person to whom one speaks, the food in front of one—these things change all the time. This means that one cannot use a notion tied to one of these relations to accumulate information about a given object. The permanent file should be one that contains objective information, that will get at the same object even though one's relations to it change. But one is always identical with the same person. My "I-buffer" can be my permanent file for myself.

5 V

If one did not move from city to city, one could also use one's "here-buffer" as a permanent file for the city in which one happened to live. Suppose one not only did not move from city to city, but did not have anything to do with other cities. One never acquired information about them, or performed actions whose success depended on them. Imagine a child, who has no knowledge that there are other cities. Such a child would not even need a here-buffer. She needs no notions of cities (or "places one lives") at all.

We often talk about the weather, as if rain and snow and sleet were states of times, rather than states of places at times. We say, "It is snowing," rather than "It is snowing here". In effect, we handle a 2-ary phenomenon with a one-place predicate. This works fine so long as we can take it for granted that we are all talking about the weather in the place where we are talking. We can also think about the weather in this way. So long

as the place in which we pick up information about the weather and the place to which we apply that information are fixed as the same by factors outside of thought, we don't need to keep track. We can have a one-place idea for a 2-ary phenomenon.⁷

Consider, for example, the way we think about time before we learn about time-zones. When we look at the kitchen clock or our watches, we learn about what time it is in the time zone we are in. Looking at one's watch is a normally time-zone-one-is-in informative way of learning the time of day. If one has very long arms and very acute vision, or lived in some city where the time zones change, one might have to be careful. But it's pretty secure. Most of the actions we use our kitchen clocks and watches to guide, are normally time-zone-one-is-in dependent ways of acting. This is because many of the things we do, such as getting up, eating, going to school, going to bed and the like, are things that people try to do when it is a certain time of day in the time zone they are in. As long as our dealings with time amount to using information gained in normally time-zone-one-is-in informative ways to guide actions that are time-zone-one-is-in dependent for their success, we have no reason to even be aware of the fact that the time of day is relative to time-zones.

We need to be aware of this relativity when this condition is broken. A child may learn how things can go wrong when she speaks to her grandmother long distance, or takes a trip across a time-zone. Her watch still tells time accurately enough, but it is the time at home, not the time at the place the child is visiting. But it is the time in the place the child is visiting that determines when lunch is served, and when the good television programs come on, and when one is expected to go to bed.

Similar remarks apply to self-notions. For many purposes we don't need notions of ourselves at all. Consider the simple act of seeing a glass of water in front of one and drinking from it. The perceptual state corresponds to a relation between an agent and a glass of water. It is the state an agent is typically in when there is a glass of water in front of that agent. The perceptual state is then not only normally object-in-front-of-one informative, but one-who-is-in-the-state informative. The coordinated motion of hand, arm and lips by which the agent gets a drink is not only normally object-in-front-of-one effecting, but also agent-who-does-the-action effecting. The identity between the perceiver and the agent is (normally) guaranteed outside of thought, by the "architectural" relations between the eyes and arms. One need not keep track of it in thought.⁸

Another somewhat Tractarian⁹ or Carnapian¹⁰ way of making this point, is to say that the world as we perceive it does not include ourselves, but has ourselves as sort of a point of origin. Suppose I tell you that one point is at (4,5) and another point is at (5,4). As long as you can assume that the points have been given, relative to the

⁷See my "Thought Without Representation," in Proceedings of the Aristotelian Society, Supplementary Volume 60 (1986): 263-83.

⁸See Dan Dennett, "Where Am I?" in *Brainstorms* (Cambridge: MIT-Bradford, 1978).

⁹Ludwig Wittgenstein *Tractatus Logico-Philosophicus* (London: Routledge & Kegan Paul, 1922) secs. 5.62ff. See also Moore's description of Wittgenstein's later views in G.E. Moore, *Philosophical Papers* (New York: Collier Books, 1962), pp. 302-03.

¹⁰Rudolf Carnap, *The Logical Structure of the World* (Berkeley: University of California Press, 1967), sec. 163.

same point of origin, you know that to get from the first to the second you take one step away from the y-axis and one step towards the x-axis. But if they are given relative to different origins, you will not know what the relation between them is. Similarly, if I show you what things look through a certain pair of eyes, you will know what an arm, connected to the body to which the eyes belong, in the normal way, will have to do to intersect with a certain cup. But you would have no idea how any arbitrary arm might have to move to perform that operation on the seen cup.

The self really comes in twice over when one notes that one is hungry or that one's hands are dirty, as both the perceiver and the object perceived. And when one eats or washes one's hands, one is both the agent and the object effected. The success of sticking one's hands under the faucet, as a response to the sight of one's own dirty hands, depends on a number of identities that are usually architecturally guaranteed. When one sees dirty hands in a certain way, it is the perceiver's hands that are dirty. When one washes hands in a certain way, it is the agent's hands that get clean. And when a perception of the first sort causes an action of the second sort in a more or less direct way, the subject of the perception is the agent of the action. We don't really need a self-notion to handle any of this. We will need one when we start to get information about ourselves in ways that are not normally self-informative.

In a world like ours, with mirrors, newspapers, lists of people who are supposed to be various places and the like, we all have many ways of knowing about ourselves that are not normally self-informative. They are just the same ways we have for knowing about others. I can look at my ticket and see when I am to leave; you can look at my ticket and see when I am to leave. The ticket gives the same information in the same way to anyone who looks at it: John Perry is to leave at a certain time on a certain day. I need an objective notion of myself to pick up the information, and a self-notion to put it to use. Unless I already have or acquire of a notion of John Perry as John Perry, I won't have any place to store the information I get from the ticket. Unless this is or is linked to a self-notion, I won't end up performing the normally self-effecting actions that I need to perform (like getting out of bed) in order to get to the airport on time.

6 VI

When he looks to the far end of the bus, Mach gets information about himself in a way that is not normally self-informative, but normally "person-I-am-looking-at" informative.¹¹ So this information doesn't pass into his self-notions; it is not combined with information gotten in normally self-informative ways. And it doesn't motivate normally self-dependent actions.

Suppose Mach looks down at his own vest and sees a big piece of lint. (Mach himself provides us with a picture of the way one's front characteristically looks to oneself.) He would have associated the idea of having a large piece of lint on one's

¹¹In many situations, looking into a mirror is a normally self-informative way of getting information. In Mach's case the mirror was far away and not set up for self-viewing, and he didn't seem to realize at first that he was looking into a mirror.

vest with his self-notion. That's what I mean by saying that the self-notion is the repository of normally self-informative perception. Now if Mach had desired not to have large pieces of lint on himself, he would have reached out and removed in a way that works when the piece of lint is on one's own vest. If he has this desire, and the idea of having lint on one's vest is associated with his self-notion, we would expect him to take such a normally self-directed and self-dependent action. That's what I mean by saying that the self-notion is the motivator of normally self-dependent actions.

But when Mach sees a piece of lint on the vest of the person in the mirror he does not act in this way. The information is not gotten in the normally self-informative way. So it is not combined with the other information in the self-notion, and doesn't lead to the action that works to remove lint from oneself.

At the beginning of the episode, Mach formed a notion for the person he saw, whom he took to be getting on the other end of the bus. This was a notion of himself, but not a self-notion. We assume Mach knew who he was, and so that he had a notion of Ernst Mach as having all of the well-known properties of Ernst Mach that was also a self-notion. But even if Mach had been in the middle of a bout of amnesia he would have had at least a self-buffer, a notion tied to normally self-informative action and perception, though not linked to any non-self-relational notion. Mach's beliefs change, during the episode, in that he comes to link the new notion formed when he got on the bus with the old self-notion or notions that he has. If, after he has made the connection, he notices that the person in the mirror has a piece of lint on his vest, he will pick the lint of his own vest in the normally self-dependent and self-effecting way of picking lint off one's vest.¹²

¹²Various versions of the material on which this paper is based have been given in lectures at Stanford University, Notre Dame University, Cornell University, the University of Wisconsin at Madison, the University of California at Davis, Princeton University, and Santa Clara University. I am grateful for the many helpful comments that have been made on each of these occasions; I'd particularly like to thank David Copp and Carol White.