

On Knowing One's Self

John Perry

Philosophy Department,
University of California Riverside

[~10,700 words not counting reference list]

In Shaun Gallagher (ed.), *The Oxford Handbook of the Self*. Oxford: Oxford University Press, 2011: 372—393

§1. INTRODUCTION

I believe that I live in Palo Alto---a rather dull case of self-belief.¹ But what do I believe, when I believe *that*? It seems the content of my belief should be just what has to be the case, for it to be true. My belief is true, if and only if a certain person, John Perry, lives in Palo Alto. So that must be what I believe. If so, it is the same thing I believe when I believe that John Perry lives in Palo Alto. But are these beliefs really the same?

It seems I might believe that John Perry lives in Palo Alto, without believing that I do. In a few years, when I am even more senile than I am, I may wander off, forgetting my name and where I live. Finding his driver's license with a Palo Alto address, I may be quite sure that John Perry lives in Palo Alto. But that really isn't *self*-knowledge, only knowledge about the

¹ `Self-knowledge' is a common and rather mellifluous phrase, usually thought to stand for a virtue; none of this can be said for `self-belief'. But this essay is mainly about self-belief; not much attention is paid to what makes some of it knowledge.

person I happen to be. I might wonder if I live there too, but think it doubtful. It seems that what I would be doubtful of, that I live in Palo Alto, is one thing and what I would be sure of, that John Perry lives there, is another.

In this case, would it be correct to say “John Perry believes that he lives in Palo Alto”? According to Hector Neri-Castañeda, it depends on whether ‘he’ is functioning simply as an anaphoric pronoun, picking up the reference of the subject term, or instead as what he calls a ‘quasi-indicator’ [Castañeda 1967/1994]. If the first, the report would be true. But if ‘he’ is functioning as a quasi-indicator, the report would be false. It would require something more of the belief, than simply that it about the subject. It would describe the sort of case where the subject would use the first person to express the belief, and so it would not describe the case we are imagining.

There appears to be more to self-belief than having a belief about the person one happens to be. I can believe, or imagine, or hypothesize, that John Perry lives in Palo Alto without believing, imagining or hypothesizing that I live there. What is the difference?

§2. FREGE

We seem to have here a special case of Frege’s problem about identity [1892]. Where A and B are different singular terms, ‘‘A = A’’ seems to have quite different “cognitive significance” than “A = B” even if “A=B” is true. The problem is not restricted to identity statements. Given the identity, any two statements “A is so and so” and “B is so and so” will predicate the same

condition of the same object. But they may not have the same cognitive significance.

The experiences that lead to the beliefs expressed by the statements may differ. I may learn that Tully was a Roman senator by reading texts that contain the name 'Tully', without learning that Cicero was a Roman senator. Even if I believe that Tully was a Roman senator and I also believe that Cicero was a Roman senator, the beliefs themselves may be different. I can acquire them at different times, from different sources. They may be altered or extinguished independently. And the different beliefs motivate different actions, most clearly what I infer and what I say. In the imagined situation, my belief that Tully was a Roman senator explains why I write 'true' next to "*True or False: Tully was a Roman orator*" on my mid-term, inferring from my belief that he was a senator, that he was probably an orator. In all of these ways, the statement "John Perry lives in Palo Alto" might have a different cognitive significance for the senile me than "I live in Palo Alto."

When late in Frege's career his attention turned to indexicals and particularly the word 'I', he said, "Now everyone is presented to himself in a particular and primitive way, in which he is presented to no one else." ([1918/1967]: 25-26). One way to read Frege's remark is to suppose that for each person there is a primitive property ψ , both uniquely instantiated by that person and unobservable to anyone else, and by 'I' the person means 'the ψ '. Something like this reading may be required to make Frege's treatment of 'I' fit into his general theory of sense, modes of presentation, and reference. But it isn't very plausible to suppose that for each person there is such a private

property that only that person instantiates, and that he associates with 'I' ([Perry, 2000]: ch. 1) .

The remark is more plausible if we reverse the quantifiers: there is particular and primitive a way of being presented to a person, in which each person is presented to himself but to no one else. Or perhaps: there are methods of knowing facts about a person that are available to every person, and are, for any given person, methods of picking up information about that very person, but not ways for that person to learn about anyone else. This is seems quite plausible. There is a method for knowing whether one has a headache, the same method for each of us, that won't work to find out if anyone else has a headache. There is a way of finding out whether one is hungry, that won't work to find out if anyone else is hungry. There is a way to find out if one needs to go to the bathroom, that won't work to find out if anyone else needs to go to the bathroom. There are ways of finding out what we feel, and believe, and worry about that we each can use to find out about ourselves, but not about others. These are not always more reliable than other methods. Parents are often better at knowing when their two-years olds need to go the bathroom than the two-years olds. There may be vexed issues about what is going on when we apply these methods, and the word 'methods' may have misleading connotations, but something along these lines seems beyond doubt. We all use these same methods to find out things about ourselves, but can't use them to find out about others. I'll call these *self-informative* ways or methods of knowing. So we might interpret Frege's idea as coming to this: we are each uniquely presented to ourselves and to no one else *as* the person we find out about with self-informative methods.

Still it seems a bit odd to say in many such cases that *we* are *presented* to ourselves at all. As Wittgenstein notes in *The Blue and Brown Books*:

...there is no question of recognizing a person when I say I have a tooth-ache. To ask, 'are you sure it is *you* who have pains?' would be nonsensical. (New York, 1958) 66-67.

When you see that someone *else* has a toothache that person is presented to you; you see a certain person, perhaps and rubbing their gums and complaining, "My tooth hurts." The person is presented to you in perception; the toothache is not presented but inferred. What is characteristic about one's own case is that the owner isn't really presented at all --- just the toothache.

3. IMMUNITY

In "Self-Reference and Self-Awareness" Sydney Shoemaker begins by considering Wittgenstein's remark. He develops Wittgenstein's point with the concept of *immunity to error through misidentification relative to the first person pronoun*. [1968/1994]. If, while wrestling, I perceive a bloody foot and think "I am bleeding," I may be wrong, because although someone is bleeding, the someone is not me. In contrast, if I see a canary, and, on the basis of visual experience say "I see a canary," it cannot turn out that I was right that someone was seeing a canary but wrong that it was me doing the seeing. In the bleeding case we don't have first person immunity, in the seeing case we do.

Shoemaker and Anscombe [1975/1994] both suggested the problem we are considering might provide philosophers grounds for thinking that

when a person thinks or asserts "I am so and so," the word 'I' doesn't refer to the person himself or herself, but a closely related object, a Cartesian Ego --- a conclusion they both reject. Anscombe, but not Shoemaker, thinks the only way to avoid this conclusion is to give up the whole idea we refer with 'I'. Like Shoemaker, she holds that in cases where we have first person immunity, we are aware of an occurrence of some property F in a special way that does not require identification or recognition of the possessor of the property, and entitles us to say, "I am F". If 'I' refers, then we will have referred to ourselves. But she maintains that we should not say that 'I' refers at all. Her thinking echoes Lichtenberg's criticism of Descartes:

We should say *it thinks*, just as we say *it lightens*. To say cogito is already to say too much as soon as we translate it I think. To assume, to postulate the I is a practical requirement.' ([Lichtenberg 2000]: Notebook K #18, p. 190.)

If I have a toothache, not in the throes of Cartesian doubt, but while communicating with others, it seems worthwhile to note, as a "practical requirement," that I know of the toothache directly, not in the way that I learn of the toothaches of others. Perhaps this is all that 'I' really does; it connotes an absence rather than referring to a presence. If we take this line, Anscombe argues, we can avoid an argument for Cartesian egos based on the special features of the first person pronoun.

Anscombe reconstructs Cartesians as being impressed by two things about the first person. The first is that when I say or think "I," I can be sure that what I refer to exists, but I may not be in a position to know that any

person in the ordinary sense, a flesh and blood being, exists. Descartes was sure that when he thought "I exist" that he was referring to something, although he wasn't sure, at least according to Anscombe, that the person Descartes existed. The second is that when I say "I", I am immune to referring to some being other than the one I intended to refer to. The seductive idea is that Cartesian egos are the sorts of things with respect to which one wouldn't need to have either of these worries. Since we don't have them when we say "I", we must be referring to Egos. Anscombe thinks that the only way to avoid the Cartesian conclusion is to give up the idea that 'I' refers, but that seems like a rather heavy price to pay; it is, as Evans says, "an extraordinary conclusion" ([Evans, 1982/1994]: 190n.).

But Cartesian Egos wouldn't solve our problem about 'I'. Suppose we each have a Cartesian ego, and when each person uses 'I' that use refers to his or her own Cartesian ego. Still, it seems that if there are Cartesian egos, they can be named. Call mine 'JP'. If one can forget one's own name, one can also forget the name of one's Cartesian ego. If I find a copy of this paper, after turning senile, I may read it and learn that there once was a fellow named 'John Perry', the author of the paper, who named his Cartesian Ego 'JP'. If my use of 'I' refers to my Cartesian Ego, then my statements "I am JP", "I am I" and "JP is JP" would all be true. But I might be as certain as that "I am I" and "JP is JP" are true, but be quite convinced that "I am JP" is false. The problem about the difference between "I am I" and "I am JP" isn't affected by whether we take ourselves to be referring with 'I' and 'JP' to Cartesian egos, or to quite ordinary persons. The problem appears to be

with the difference in ways of referring, not in the nature of what is referred to.

In one clear sense, I *can* be wrong about the reference of my use of 'I', whether 'I' refers to me or to my ego. Suppose that in my senility, I not only forget that I am John Perry, but come to believe that I am John Searle: years of admiration morph into delusion. I may then think I am referring to John Searle when I say, "I live in Berkeley": I think that I am referring to John Searle *by* referring to myself. The move to egos makes no difference. Call John Searle's Cartesian Ego 'JS'. I may think that I am referring to JS *by* referring, with 'I', to my own Cartesian ego. Since Cartesian egos do not eliminate our problems, our problems do not motivate the postulations of Cartesian egos.

Anscombe thinks that our paradigms of reference, proper names and demonstratives, do not provide the guarantee of existence and the immunity from misidentification that 'I' does, and so that the whole idea that 'I' refers is fishy. Shoemaker argues more plausibly that we simply need to appreciate that 'I' refers in a different way, putting different cognitive demands on the user, than do descriptions, names or even demonstratives. With 'I' there is no issue of identifying the thing designated as meeting certain further criteria.

With demonstratives there is a sort of immunity. I might refer to a cup I see and hold in my hand as "this cup". If I am not thinking of it as some cup I have previously encountered, or the cup that someone else has asked about, but simply as the cup I see and hold at that moment, there is no issue of being wrong about *which* cup I refer to. Still, there may be other

candidate cups around. The words 'this' or 'that' permit us to refer to different salient cups, depending on our intention. But if I am using 'I' with its ordinary meaning, there is no additional intention that is needed or can make any difference: I refer to myself.²

Evans notes that on these issues there is a closer fit with other indexicals --- 'here' for example, and especially 'now' and 'today'--- than with descriptions, names or demonstratives [Evans 1982/1992]. Just as there is immunity to misidentification of persons relative to 'I' there is also immunity to misidentification of times relative to the word 'now' or 'today'. Suppose the time of my hunger pangs and of my utterance is July 4. Still, "I am hungry *on July 4*" is not immune to error in the way that "I am *now* hungry" and "I am hungry *today*" are. I *might* be wrong about what day it was, and think that my hunger was occurring on July 4 when it was really occurring on July 3. The phenomenon of hunger presents itself in a certain way; when it does so, I can be sure that the 'owner' of the hunger is *me*, and the time of the hunger is *now*; but neither the owner nor the time seem to present themselves as further objects.

This suggests that the nature of the immunity that attaches to 'I', 'now' and 'today' has to do with what they have in common in the *way* they refer.

Indexicals differ from names and descriptions in that they are associated, by their meanings, with what I shall call *utterance-relative roles*. The object that occupies the role, relative to a given utterance of an indexical, is the object to which the indexical refers; unlike the case of demonstratives,

² For a discussion of the role of intentions with demonstratives; see [Perry, 2009a].

a further intention, beyond using language in the normal way, is not required.³ An utterance of 'I' refers to the *speaker* of the utterance; an utterance of 'now' refers to the *time* of the utterance; an utterance of 'here' designates the *place* of the utterance. Indexicals are assigned to roles, because the rule of reference that governs them is *reflexive*. An utterance of the word 'I' refers to the speaker of *that very utterance*. An utterance of the word 'now' refers to the time of *that very utterance*. In contrast, the rule of reference for a name is not relative to the utterance, at least on most theories. Rather than an utterance-relative role for the referent to play, a property for the referent to have, independently of any particular utterance of it, seems assigned to the name as its meaning. On a popular theory, the rule of reference is something like this: an utterance of the name 'N' refers to the object to which 'N' has been assigned by the convention. Descriptions, at least those that do not themselves include indexicals or demonstratives as components, are also not utterance-reflexive: an utterance of 'the ψ ' denotes x , if x is the unique object that has the property assigned by the lexical and compositional rules of the language to ψ .

Can the utterance-reflexivity of 'I' shed light on our problem?

Anscombe thought not:

...we are inclined to think that 'It's the word each one uses in speaking of himself' explains what 'I' names, or explains 'I' as a "referring expression". It cannot do so if 'He speaks of himself' is compatible

³ I actually believe things are somewhat more complicated, so that intentions can play some role with respect to 'now' and 'here', in determining the extent of time or space referred to. If we reflect on time-zones and long-distance calls, we can even find role for intentions to play with 'today'. See [Perry, 2009b]: chapter 4.

with ignorance, and we are using the reflexive pronoun, in both cases, in the ordinary way. (142)

To use the reflexive pronoun 'himself' in the ordinary way, is to mean no more than this:

[I]: For each x , when x uses 'I', x refers to x .

The other way of using 'himself' is as a quasi-indicator. Then the rule requires more than that the user and the person referred to are one and the same, more than [I] does. It would require that the user is thinking of himself in the way appropriate for using 'I'. But so conceived, the rule only gets at the meaning of 'I' in a circular way.

Pace Anscombe, I do think this simple rule [I], and in particular its reflexive nature, does explain what is special about the first person, why it is apt for expressing self-beliefs, and provides a clue to the difference between self-beliefs and beliefs merely about the person one happens to be. The key point is that the rule assigns a role to the word, and roles play a central role in thought and action that self-reference and self-belief require us to recognize.

§4. ROLES

Consider 'today'. The rule of reference for 'today' is more or less this: An utterance of 'today' refers to the day on which the utterance was made. The conditions for being the referent of an utterance of 'today' are 'utterance-bound'. The day has to stand in a certain relation to the utterance, to play a certain role vis-à-vis the utterance. In contrast, the reference-conditions provided by 'July 19, 2009' are not utterance-bound. Whether an utterance of

John Perry 5/12/10 10:23 PM

Deleted:

those words occurs before that date, or on that date, or after that date makes no difference. Whenever the utterance occurs, it refers to the 19th day of the seventh month of the ninth year of the third millennium.

Now consider the methods we have for finding out about the current day --- that is the day we will refer if we say, 'today'. "If you want to know whether the sun is shining on a given day, look into the sky. If you want to know whether it is raining, look out the window." These methods are not *utterance-bound*, for one can learn a lot using them without uttering anything. But they are *inquiry-bound*. They are methods for knowing about the day on which the inquiry is conducted. Adapting a slightly modified version of Frege's formula, we can say that there is a way information about each day is presented to us, on that day, a way in which information about that day never was presented to us before, and never will be again.

Other methods of inquiry are not inquiry-bound, or, more carefully, less inquiry-bound. Once July 19th is in the past; one can no longer find out about it by looking at the sky or out the window. But you can look at your diary, or go the library and find the *New York Times* for that day and find out a lot about it.

I will say that a method of reference, or a method of finding out about an object, is *role-based* if the object referred to, or the object about which information is obtained by using that method, is one that plays a certain role vis-a-vis the act of reference or act of inquiry. I'll say that such methods are *property-based*, if the object is determined by some property that it has or lacks independently of the act of reference or inquiry. So, at least arguably,

proper names, definite descriptions that are free of indexicals and demonstratives, and dates give us methods reference that are property-based rather than role-based.⁴

Role-based methods of reference and inquiry belong to a wider class of role-based ways of acting. When we move our limbs or body so as to have a direct effect on an object, *which* object we have an effect on will be determined by the relation it stands in to us. There is a certain way of picking up a cup of coffee, that is a way of picking up the cup of coffee in front of us. It will work for any person, and any cup of coffee, as long as that cup of coffee is in the right relation to --- plays the right role in the life of --- the person at the time of action. And of course there are ways of finding out about the cup that similarly bound to the act of inquiry. By opening your eyes and looking in front of you, you can find out more about a cup that is in front of you. By sticking your finger in it, you can find out if the coffee in it is hot. By bringing it to your lips and tilting it, you can find out how the coffee in it tastes.

This example illustrates that many of the roles that objects play in our lives at a certain time are *architecturally bound*. Facts about how the world works, about how we are built, and about how things are built and arranged, insure that the object that plays one role in our lives at a certain time will also play many others. The cup of coffee I can see by directing my glance in a

⁴ Again, I think things are actually a bit more complicated. Ultimately all activity including all thought and all use of language connects to the world in role-based ways, but the roles can be of quite different sorts, and less and less dependent on the features of the particular episode and more and more on external features that are the same at different times and places and for different people, and thus approximate to being property-based. See [Perry, 2009b].

certain way, is the one I can pick up by moving my hands and arms in a certain way.

If you live a temporally simple life --- live day to day, as we say --- never making appointments for more than a few hours ahead of time, having no calendars, basing what you do each day on what it is like and what happens that day--- you really don't seem to need the term 'today'. You don't need to think, "It's cloudy today." You can just think, 'it's cloudy'. You don't need to think, "I have an appointment today at 4 p.m.," you can just think "I have an appointment at 4 p.m." If you think "it's cloudy" on a certain day, your thought will be true or false based on what it's like on that day. But that dependence needn't be reflected in an articulated component of your thought. One can imagine simple folk who find out about each day in the way we do, and base what they do on what they find out, but who don't really think about *days* at all. Of such folk, I will say that their thoughts and language *concern* the present day, the day on which their inquiries and discourse and other actions occur, even though there is no word in their vocabulary or idea in their minds specifically charged with *standing for* the present day. That the truth-conditions of their thoughts and discourse have to do with the day on which they occur, is determined by the fact that the methods of inquiry they use provide information about that day, and the success of the actions these thoughts motivate are determined by how things are on that day. It is their *form of life* that ties their thought and language to the present day, rather than any word like 'today' (see [Perry 2000]: ch. 10). If one found such folks using the word 'today', one might wonder if it really

referred, or simply served redundantly to register the methods they were using to find out about cloudiness, rain, and the like.

§5. THE SELF

Being the present day is an example of an important role. Such roles are architecturally bound with others that permit one to get information about the objects that play them in certain ways, and determine how such objects affect the results of one's action. The present day is the one I refer to with 'today,' the day I can find out about by looking out the window, the day whose weather will determine whether lying on the beach will make me tan or burnt or cold or wet. I'll say that such roles 'support' role-based methods of inquiry, and role-sensitive ways of accomplishing things. This latter phrase includes the fact that we have ways of acting that affect the objects that play the role, and ways of acting whose success is dependent on the attributes of the objects that play the role, although not effecting them in the ordinary sense. If I say 'The coffee in that cup is cold,' the truth or falsity of my statement will be sensitive to the cup I refer to, although I don't really affect the cup by referring to it and saying something about it.

Selves are persons; *Self* is a role persons play. One's self is like today or one's home or one's father; a perfectly ordinary object, thought of in virtue of its relation to the thinking agent. The relation associated with 'self' is just identity. That is,

If x is y 's self, then $x = y$

There are self-based methods of inquiry, and self-sensitive ways of accomplishing things. Reflection and introspection and perception are all

self-based methods of inquiry. Perception is doubly role-based. I find out about the objects seen and touched. And I find out about the see-er and toucher. All basic ways of acting are similarly doubly role-based. I have an effect on objects that play a certain role in my life at that time: the cup I see and pick up; the telephone I talk into; the person at the other end of the line connected to the phone I talk into, and the further results of my action depend on their properties. But the actions are also ultimately self-sensitive; the effect of one's actions depends on how things stand in relation to oneself. A certain movement will result in *my* head being scratched: it affects me. Another movement will result in the cup in front of *me* being moved across the table; or the phone *I* am dialing being connected to the phone whose number I dial, or the person who answers that phone being distracted from some more important task. By moving my lips in a certain way, I can tell the person on the phone with *me* that I'm sorry for dialing the wrong number; I don't need to know anything about who they are except that they stand in this relation to me.

Philosophers, worried about the metaphysics and epistemology of consciousness and the specter of skepticism, and charged with recruiting future generations of philosophers by working yearly through the *Meditations* with undergraduates, focus on self-informative methods, like introspection, that seem to be *necessarily* self-informative. There is a much wider class of normally self-informative methods, architecturally and contingently guaranteed to be self-informative. There is a way of finding out about your own feet, looking straight down and inspecting the feet you see, but distorting glasses, or very bad posture, could make them methods of

finding out about the feet of the person next to you. Looking in a small mirror held close is a way of finding out about one's own nose, lips, and beard, but if we had large donut-shaped heads it might be a way of finding out about what was happening behind us. Ordinarily the metaphysically guaranteed and the merely architecturally guaranteed self-informative methods function as a seamless whole; we think of ourselves as the occupants of a bundle of roles bound together.

We might ask what justifies us in thinking, say, that the being that is aware of *this* toothache is the same as that which is aware of being aware of the toothache ([McTaggart, 1921—27]: §386). H. P. Grice once proposed solving this problem by introducing the self as a logical construction; the relation of possible co-awareness by a single act of introspection is used to define a “total temporary state”; then a version of Locke's memory theory is constructed to bind the total temporary states into a single self or person [Grice, 1941/2008]. But, as with most attempted logical constructions, there is a heavy dependence on counterfactuals, that seem to rely on causation, and ultimately suggest that the unity that holds the various aspects of our selves together has to do with dependence on a single continuant like a brain or a body. [Perry, 1975/2008]. At any rate, the question doesn't usually arise; I am the one who feels this toothache, and the one whose owns these hands I see before me, and the one whose feet I can move at will, and the one who is looking at the things I see. I don't have to decide which really captures my sense of self, any more than I have do decide whether the car I drive is the one I sit in, the one I steer, the one I brake, or the one whose windshield I see through.

When we investigate the objects around us, in virtue of applying role-based methods of inquiry, we pick up information that motivates role-sensitive methods of acting. I learn that a the cup I see contains hot coffee by employing role-based methods of inquiry; I pick it up and drink from it by using role-based methods for picking up a cup and drinking from it. But there being a cup of hot coffee there isn't sufficient motivation for so acting; the act would be irrational unless I was thirsty, or needed caffeine, or both, and unless I knew that the cup was in front of *me*. So my beliefs need to reflect those facts about me, for my act to be properly motivated.

In this case, it suffices that I have beliefs that *concern* me. My beliefs can be like Lichtenberg and Wittgenstein and Anscombe have in mind, of the sort we might express with "Hot coffee in that cup in front; there is thirst and caffeine deprivation; so drink." We can imagine animals, cognitively sophisticated enough to perceive the world in terms of objects having properties and standing in relations, and perhaps even to re-identify objects perceived at different times, with no need to appreciate *themselves* as objects. Their beliefs concern them, but do not represent them in the way that they represent other objects.

Such an animal picks up and acts on the basis of information about itself in spite of not having an idea that stands for themselves --- much less a first person pronoun. It gets information about how things are around *it*, and this influences which self-sensitive actions it takes. I'll say such an animal has *primitive* self-knowledge, gained by methods that are self-informative, and motivating actions that are self-sensitive. And we are like such animals when we are young enough, and revert to this more primitive level of thought

in cognitively undemanding situations (and, oddly enough, in the rather more demanding situation of trying to follow Descartes in methodical doubt.)

Lichtenberg's insight is that in the situation of the Second Meditation, where Descartes is (we are to suppose) not only meditating, but withholding assent from the very idea that there is anyone else to communicate with, he has no real need for 'I'.

The use of the particle 'self' in "self-informative methods," may give the impression of circularity. I can learn that I have mustard in my beard in a self-informative way, by looking in a small mirror held right in front of my face. Or I can learn in a way that others can use --- perhaps my face and mustard-soaked beard appear on the Jumbo-Tron above the outfield at a Giants baseball game, and I see it, and learn about myself, in the same way that everyone else learns about me. In both cases John Perry learns something about John Perry---about *himself*, in what Ancombe calls the "ordinary" --- that is, not quasi-indicative --- use of the reflexive pronoun. The difference is that with the first method of knowing, but not the second, the known is always, or at least normally, the knower. Marking the difference doesn't require any appeal to an antecedent or underlying self, or any use of quasi-indication, and is not circular.

§6. THE PERSON ONE HAPPENS TO BE

Most of us, even if we would like to live only in the present day, can't manage it. In addition to our role-based ways of thinking and talking about days, we have an elaborate property-based ways of thinking about them. Calendars embody this. I have a large calendar on my refrigerator, where I

make notes about appointments and projects. In the 'July 21' box, there is a note, "John has dental appointment". The notation in the box amounts to a statement, that is true if I go to dentist on July 21st, false if I don't. It expresses the same proposition on May 1st, when the appointment was made and noted on the calendar, as it does on July 21st, when I will or won't go, and as it will when that date is long past.

Our system of dates gives us a property-based way of referring to and thinking about days that is independent of when we are doing the referring and thinking. It is useful because the world we live in provides us ways of forming well-grounded beliefs about days other than the present one. And the world also provides us with structures so that our actions can have predictable results on days other than when we perform them. When the dentist's secretary and I agree on the July 21st appointment, we come to have well-grounded beliefs about July 21st: that I will show up, that the dentist will be ready to look at my teeth, and the like. When, on May 1st, I write a note about the appointment in the July 21st box on my calendar, my action, if things go right, will have its intended effect on July 21st, when seeing the note will lead me to drive to the dentist.

The calendar doesn't have any boxes marked "today". But I have a little red magnet that I move each morning to the appropriate box. Something like this happens automatically with calendars on computers. What is the point of that? What information will I add to the 'July 21st' box on the calendar when I place the magnet there that morning? It is not the magnet, but the fact that the magnet is there on a certain box on a certain day, that carries the information. If my habits are regular enough, the magnet's

presence in a box on the calendar will carry the information that the day *when* I am doing the looking is the very day the box stands for. The box I see the magnet on, has notations from which I can discover what is supposed to happen on the day the box stands for, and the magnet tells me I can find out more about that by looking out the window. It is the day on which I will do the things, I do today. I know how to go to the dentist --- get in the car and drive up Channing to his office --- but the method I use is a method for getting to the dentist on the whatever day that I use it. There is no specific method for getting to the dentist on July 21st. To do that, I need to deploy the role sensitive method, on the day when my little magnet is on the same box as ``John has an appointment at the dentist''; that is when July 21st occupies the `today' role.

To navigate around the more complex informational world that dates and calendars enable to me to live in, I do need the word `today' and the thoughts I can think with it. I need to be able to make inferences like: "I go to the dentist July 21st (what the calendar tells me); today is July 21st (what the magnet tells me); so I go to the dentist today". All the calendars in the world won't make me into a well-organized and punctual person, if I can't convert the property-based references provided by the dates, into my inquiry- and action-bound ways of thinking of the present day.

The world also supplies property-based ways of referring to me and thinking about me, most prominently my name. I'll call these *nominal* ways of referring and thinking. This in turn makes possible ways of gaining information about me, and acting in ways that affect me, or whose success is sensitive to my facts about me, in ways that are not self-sensitive. If

someone knows my name, she can look up my phone number in the directory, find out how much I owe on my property taxes at the County Hall of Records, and, in the age of the internet, no doubt find out endless number of other things about me by googling my name (although this is mostly a way of finding out about John Perry Barlow, the Grateful Dead lyricist and cyberlibertarian). Although the County is nice enough to send me a bill, once I open it I use the same method as anyone else to find out what my property taxes are: I look for a notation where the name 'John Perry' is linked with an amount due. If I forget my own phone number, I will find out what it is in the same way anyone else would, by looking it up in the directory. And each quarter I look in the class schedule to find out when and where my class meets, the just as the students do.

When I am senile, and forget my own name, I will still be able find out about myself by reading about John Perry in various directories and schedules, and hearing his name in conversation. This won't be what we ordinarily think of as self-knowledge. What is lacking is the ability to get the knowledge thus gained tied to knowledge of myself gained in self-informative ways. Someone who was stuck in an Lichtenbergian language, with no first person way of referring and thinking, would be similarly disabled. To make use of these public, non-self-informative methods of finding out about people, in order to gain self-knowledge, I need a way of thinking about and referring to myself that is associated with self-informative ways of knowing and self-sensitive ways of acting, but can also be linked to nominal methods. I need to be able to look at the time schedule, and infer: "John Perry teaches in the Philosophy Department Seminar Room at two

p.m.; I am John Perry; so I teach in the Philosophy Department Seminar Room at two p.m.”

So, perhaps surprisingly, it is the availability of methods of gaining information about ourselves in non-self-informative ways that makes having a first person crucial. In an informationally simpler world, where there are no such methods, we wouldn't need a first person any more than those who live in the present moment, or at least in the present day, would need 'today'.

We need to be careful here, however. It is not really the first person that is needed. It is the function that the first person performs that is important. A name could perform it, and often does with children ([Perry 2000]: ch. 3). If Eloise is the only child with that name in her household, so that every remark she hears that uses the name 'Eloise' is about her, and every time she hears a shout 'Eloise!' it is she that is summoned or chastised, then hearing the name 'Eloise' is a self-informative method of gaining information for her, and she can use 'Eloise' in asking for things in much the way she will use 'I' later on: "Eloise wants a cookie." The word 'I' is suited by its meaning to perform this function for every English speaker, and Eloise will need to learn it to get cookies when she starts interacting with other Eloises, and people who don't know her name.

§7. IDEAS, NOTIONS AND BUFFERS

It may be helpful to develop an admittedly over simple model of cognition that organizes the ideas discussed so far. I'll assume we have *beliefs*, which are particular states. Beliefs have causal roles and contents, which *mesh*. Our ordinary way of identifying beliefs is by their contents, for example, my

present belief *that my shirt is badly coffee-stained*. The content is a proposition, that my shirt is badly coffee-stained; a proposition is an abstract object that encodes truth-conditions. How we think of propositions, as sets of possibilities, or structures of objects and properties, won't matter much so long as they encode truth-conditions in one way or another.

The belief I have about my shirt was caused by a self-informative perception. I look down and see my shirt with some big coffee-colored splotches on it. My perception has content, that the shirt had coffee-colored splotches on it. My perception, plus some background beliefs and memories (of spilling coffee) led to the conclusion that my shirt is badly coffee-stained. When we say the perception led to the belief, we have in mind two things. First, the perception was a cause of my acquiring the belief. Second, the content of the belief makes sense given the content of the perception, together with the contents of background beliefs and memories. That's part of what I mean by saying the cognitive and causal roles *mesh*. We also get meshing when we look forward. My belief that my shirt is coffee-stained, together with my desire to appear reasonably well-kempt when my wife and I go out to dinner, together with a number of other beliefs, such as that I have a clean shirt upstairs in the closet, will motivate me to go upstairs and change my shirt. That is, it will cause me to execute a number of bodily movements, which will get me upstairs in front of my closet, and then result in getting a fresh shirt out of the closet and putting it on. The results these movement will have, *if* my beliefs are true, are that I will have changed my shirt, and will have a cleaner shirt on than I do now, and will appear considerably more kempt. And these results will promote the satisfaction of my desire, to

appear well-kempt when we go out for dinner, although of course they will not absolutely guarantee it, as there will be more opportunities to spill things on myself. So that is another part of the idea of meshing. The result of the actions a complex of beliefs and desire cause, should promote the satisfaction of the desires, if the beliefs are true.

My beliefs have structure; they are made up of ideas. Ideas contribute to the causal roles and contents of beliefs of which they are components. And they are common components of many beliefs. My idea of a university is a common component of my belief that Stanford is a University, my belief that U.C. Riverside is a university, and many others. My idea of Stanford is a common component of my belief that Bratman teaches at Stanford, that CSLI is located at Stanford, and many others. And so forth.

I'll just call ideas of the sort we associate with general terms, verbs, prepositions, and the like 'ideas'. I'll call ideas associated with particular things --- people, places, things, universities, and, although it's a bit of a stretch to call them 'things', days--- 'notions'. A notion, plus all of the ideas associated with it in belief, makes up a *file*. The file associated with a given notion determines the conception, or a conception, that a person has of the object the notion is *of*.

Finally, among notions, I'll distinguish between *buffers* and *standing notions*. A buffer is typically tied to a perception, or series of perceptions of an object. It is the basis of a sort of temporary file, used to keep track of the object while one stays in more or less steady perceptual contact with it.

More generally, buffers are tied to roles and bundles of roles that a given object plays in our lives for a while.

Standing notions are more or less permanent notions that we use to store information about objects we are not perceiving or otherwise epistemically attached to; objects that are not currently (as far as we know) playing any informative or action-affording roles in our lives. I have a standing nominal notion of, say, Joseph Biden, our vice-president. I've never met him; I've read about him; I've used nominal-methods to acquire information about him. Perhaps some day I will go to a Democratic Party fundraiser, and will see a distinguished, silver-haired, rather loquacious fellow, and form a new buffer. Eventually I accumulate enough information in the buffer about how this fellow looks, his opinions, his demeanor, and the like, to infer that he must be Joe Biden; that is, I recognize him; I pool the information in the buffer and the standing notion. Now I have a belief that Joe Biden, the vice-President, is across the room from me. This combines with the long standing desire to shake hands with Joe Biden, so I cross the room and so on.

What notions contribute to the contents of the beliefs of which they are components is what they stand for, what they are *of*, not the conceptions associated with the notions. The object an agent's notion is *of* need not be a good fit for his conception of it. It's nice if the ideas associated with a notion pretty much accurately depict the object the notion is of, but it is in no way guaranteed. Perhaps I wrongly believe that Biden is a former Senator from Pennsylvania. My *conception* of Biden, the complex condition determined

by my Biden file, doesn't *denote* anyone, for no one is both vice-President and a former senator from Pennsylvania.

The point of introducing this machinery is to allow us to focus on what's special about what I call 'the *self*-notion'. I assume that all normal humans past a certain age have a *self*-notion, an idea that is dedicated to dealing with information about themselves. This is the notion that controls the use of the first person; that is, beliefs of mine with my self-notion as a component I will express with 'I' or 'me', even after, in my senility, I have forgotten my name. It is the notion with which all of the information picked up in self-informative ways is associated, and the beliefs associated with this notion motivate self-sensitive actions.

Normally our self-notions will not only be associated with ideas picked up in normally self-informative ways, but all sorts of additional information, about our telephone numbers, salaries, class-schedules, doctor's appointments, and lots of other things, that we have learned about ourselves by hearing or reading our name in various letters, on various lists, announcements and the like, and hearing others say things about us. That is, for one who knows his own name, one's self-notion will also serve as a standing notion, a nominal-notion associated with one's own name. The information we get about ourselves using nominal methods and our own names, as when we look up our own phone numbers, are pooled with the information we get in normally self-informative ways, and the combined information guides our actions.

The semantics of the self-notion is very simple: it is *of* the person it belongs to. Beliefs with the self-notion as a component are *about* the person it belongs to, whether the beliefs were acquired in self-informative ways or not. From one's self-notion we can retrieve one's self-conception, what one takes oneself to be. This is quite relevant to one's psychology, but irrelevant to the semantics of one's beliefs. Psychologists tell us it is unlikely to be completely accurate. But for most of us, it is quite extensive and detailed, and gets a lot of things right. I know my own height, weight, salary, age, and the middle names of my parents. I don't think I know all of that about anyone else. But my self-notion would still be *of* me, even if it were wildly inaccurate.

Self is an extremely special role. And it has some very special features. Most importantly, it is played permanently by the same person. One's self never changes, in the sense that only one person ever occupies that role. The occupant of the self-role, unlike the occupant of the person-I-am-talking-to role, or the today-role, or the here-role, never gets replaced. This means that one can use one's self-buffer as one's standing notion for oneself. One doesn't need a little magnet to keep track of whom one refers to with 'I'.

The self-buffer isn't unique in this respect. For most of us, the same planet will play the role, planet-I'm-standing-on-or-flying-above our whole life. For many generations of people, millennia-I-live-during was a role occupied by the second millennium throughout their entire life. These roles are, or were, only contingently guaranteed to be occupied by the same thing. Universe-I-am-a-part-of seems metaphysically secure. And, as Ken Taylor once pointed out, nothing is quite so immune from error through

misidentification as the occupant of the role, possible-world-we-are-in is under the phrase, 'the actual world'.

Philosophers have sometimes tried to give simple role-based analyses of our sense of self: "the being having *this* experience," perhaps. But in fact we play many roles relative to episodes of thought, perception and action. I am the person seeing *these* things; the person having *these* sensations; the person performing *these* actions; the person experiencing *these* emotions, the person having *this* inchoate background of bodily feeling, and much else besides. Descartes maintained a sense of self, while doubting that he had a body; the special access he had to his own ideas seemed to suffice. Some schizophrenics think that their own thoughts are not their own; they nevertheless maintain enough sense of self through connections to actions and the physical world around them to think the thought that some of their thoughts are not their own. In "Where Am I," Dan Dennett prizes apart the structure of roles bound together in our ordinary sense of self, until we have no idea where we are, and only the barest sense of who we are. [Dennett 1981].

§8. SELF-BELIEF

Let's return to the problem of §1. How can I have a belief that John Perry lives in Palo Alto, without believing that I live there? What more is required for self-belief?

The current dominant view of the semantics of names and indexicals is that they are *directly referential*, in David Kaplan's sense [Kaplan 1989]. This means that they contribute the objects referred to, rather than any modes

of presentation of those objects, to the *contents* of the statements containing them. This means that the statements "I live in Palo Alto," spoken by me, and "John Perry lives in Palo Alto" have the same content, i.e., express the same proposition. If we are referentialists, then we must admit that at the level of content we lose track of the differences between the ways that indexicals and names refer, and the differences between the beliefs that might lead to using one rather than the other.

The direct reference picture does not imply that there are no special modes of reference associated with indexicals. As Kaplan notes, the rules of reference for indexicals, what he calls *characters*, do supply something very much like the modes of presentation Frege appeals to in his theory. The difference is that for the referentialist, these modes do not become constituents of the propositions expressed.⁵

The conclusion we should draw from this, however, is not that my statements, "I live in Palo Alto" and "John Perry lives in Palo Alto," have the

⁵ For Kaplan, the reference of a term is *indirect*, if the referent of a term depends on some reference-determining aspect of the proposition that statements using the term express. This is Frege's picture. The proposition expressed doesn't depend on the referent of the term; the referent of the term depends on the proposition, in particular that aspect of it contributed by the sense of the term, an identifying condition. This aspect, plus facts about the world, determines the referent. On Kaplan's picture of direct reference, the character of the term and certain aspects of the world, the context, determine the referent, and the referent determines (along with the rest of the statement), which proposition is expressed. The route from utterance through character and context to referent may be quite complex and in that sense indirect, but that's not the sense Kaplan has in mind.

same cognitive significance, but rather that, if we embrace referentialism, we should not expect to find differences in cognitive significance at the level of content. When we want to see how the truth-conditions of our statements and thoughts, the information we pick up in perception, and the results we expect to achieve in action, all *mesh* in a coherent way, we cannot lose track of the modes of reference and presentation involved. The place to look for cognitive significance is in the whole structure of the statements (and the beliefs that lead to them), the way in which the meaning of the sentence used interacts with the context to determine that content.

In particular, we need to look at what I will call the *conditions of truth*. This is a condition that the utterance or statement itself has to meet, in order to be true, that incorporates, rather than loses track of, the way reference is achieved.

Consider statements of "I live in Palo Alto," by a group of diverse people. These statements all have something important in common. Each statement is true if and only if the person who makes *it* lives in Palo Alto. Each statement *s* will have meet the same reflexive condition of truth:

If *s* is an utterance of 'I live in Palo Alto', *s* is true iff the person who makes *s* lives in Palo Alto.

On the other hand, they each have something important that is not in common, and corresponds to the referentialist theory of content. My statement will be true iff John Perry lives in Palo Alto. Debra Satz's statement will be true iff Debra Satz lives in Palo Alto. John Searle's statement will be true iff John Searle lives in Palo Alto. And so forth. Each

statement will express a different proposition, once we fix to whom the utterance of 'I' in it refers.

Now consider everyone in the group saying, "John Perry lives in Palo Alto". Assume they are all employing the same convention *C* associated with 'John Perry'. Then we can say:

If *s* is an utterance of 'John Perry lives in Palo Alto', *s* is true iff the person associated by *C* with the name 'John Perry' lives in Palo Alto.

This is not a reflexive condition of truth; the variable *s* does not occur on the right side of the 'iff'.

Unlike referential contents, conditions of truth keep track of the roles objects must play, or the properties they must have, in order to be referred to. This is required if we are to understand how the truth-conditions of the statements one makes mesh with the beliefs that lead one to make them.

Each sincere person who says, "I live in Palo Alto," will be motivated by a belief consisting of his self-notion and the idea of living in Palo Alto. Such a belief is true, iff the owner of the self-notion, the thinker of the thought, lives in Palo Alto. In normal circumstances the role of being the thinker of a thought, and being the speaker of words motivated by the thought, are tightly bound. The use of the word 'I' to refer to oneself, is just one example of a self-sensitive action; the person the agent refers to is the agent.

Each sincere person who says, "John Perry lives in Palo Alto" will be motivated by a belief consisting of a nominal notion associated with the name 'John Perry' and the idea of living in Palo Alto. In normal

circumstances (and putting aside the complexities due to multiple conventions for the same name), using a name N is a way of expressing a belief involving a nominal notion associated with N.

Until I become senile and forget my name, my self-notion will serve as a nominal notion associated with the name 'John Perry'. So, until then, my statements 'John Perry lives in Palo Alto' and 'I live in Palo Alto,' will not only have the same referential content, but will be motivated by the same belief, the association of my self-notion with the idea of living in Palo Alto. This belief won't be the total motivation, otherwise we couldn't explain when I said one, and when I said the other. How I express the belief on a given occasion will depend on what I take the conversational situation to be. If I am with some people who don't know my name, but are interested in Palo Alto, I might volunteer, 'I live in Palo Alto'. They will learn from this that they are talking to someone who ought to know something about the place, and can direct questions at him; I couldn't achieve the same result with 'John Perry lives in Palo Alto'. If I am at a Pomo bar in Berkeley and have overheard a conversation in which some mean looking people have said many negative things about analytical philosophers and Palo Alto, and am asked, 'Do you know of any analytical philosophers who live in Palo Alto,' I might say, 'John Perry lives in Palo Alto,' and leave it at that.

Suppose now I have forgotten my name, but still remember that I live in Palo Alto. I keep running into references to John Perry and believe that he lives there too. Then my statements will still have the same referential content, but not express the same belief. And, finally, when I have forgotten where I live, it will be correct to say, using the 'he' quasi-indicatively, 'John

Perry believes that John Perry lives in Palo Alto, but John Perry doesn't believe that he lives in Palo Alto." What I lack is a belief that has the referential content that John Perry lives in Palo Alto in virtue of having my self-notion as a component.

§9. SELF-NOTIONS AND CARTESIAN EGOS

Apart from being metaphysically more plausible (I hope), does the theory of self-notions fare any better with our problem than the theory of Cartesian Egos? Let 'JP' now stand not for my Cartesian Ego, but for my self-notion. Surely I might doubt that I live in Palo Alto, but be quite certain that JP's owner lives in Palo Alto. So what have we gained, as far as our original problem goes?

Self-notions are objects, and there are many ways to think about and refer to any object. Senility isn't even required for the envisaged situation. Suppose that Krista Lawlor is giving a lecture on self-knowledge at 10 a.m., and considers my theory among others. For the sake of argument, let's even assume she accepts the theory. To illustrate it she says, "Let's name Perry's self notion 'Bruno'." She writes "'Bruno is a self-notion'" on the board, where she also writes, "Bruno belongs to a Palo Alto resident." Since she accepts the theory, and knows I live in Palo Alto, she believes both statements.

At the end of class Lawlor uncharacteristically doesn't erase the board. I teach in the same room at 11, and want to talk about self-notions. I see that 'Bruno' is clearly the name of a self-notion. So I just start talking about Bruno. I refer to my own self-notion with 'Bruno', but I don't know it. I

have no idea whether the statement on the board, "Bruno belongs to a Palo Alto resident," is true or false.

This all isn't a problem for the theory of self-notions. Referring to Bruno is one thing, having Bruno as a component of one's beliefs is another. Lawlor's beliefs have her nominal notion *of* Bruno as a component; they are thus *about* my self-notion; it is part of their subject matter. Bruno is a constituent of the referential content of her belief and of the statements she wrote on the board. But it is *not* a component of her thought. It is not a notion that is a structural part of her belief, but a notion her belief is about.

The belief I express with "I live in Palo Alto" does have Bruno as a component, but it is not about Bruno; Bruno isn't part of its subject-matter. The thought is not *about* Bruno, but about me. On the other hand, the belief I express with "Bruno's owner lives in Palo Alto," like the thought Lawlor expresses with that sentence, does not have Bruno as a component; that is, it does not have my self-notion as a component, even though it is about it. It has a nominal notion *of* Bruno as a component, but not Bruno itself. If I were to assert, "Bruno's owner is no doubt a profound thinker," that would not show that I was immodest, for it would not show that I had a belief whose condition of truth is that the owner of that very belief is profound. The fact that I can think and talk *about* my self-notion without knowing that I do so is not a problem.

On the theory on offer, the differences between my senile thoughts, "I live I Palo Alto" and "John Perry lives in Palo Alto" are: (a) at the level of what goes on in my mind, that the first has my self-notion as a component,

and the second has a nominal notion of me as a component; (b) at the level of content, even though both beliefs have the same referential content, the structures whereby they come to have that content are quite different. The first has a reflexive condition of truth, and is about me and is about me in virtue of a role I play in the life of the believer, namely identity; the second has a non-reflexive condition of truth, and is about me because of a property I have, namely being the person assigned to 'John Perry' by the relevant convention.

§8. REFERENCES

- [Almog et. al. 1989] Almog, Joseph, John Perry and Howard Wettstein, eds. *Themes From Kaplan*. New York: Oxford University Press, 1989.
- [Anscombe 1975/1994] Anscombe, G.E.M. The First Person. In *Mind and Language*, edited by Samuel Guttenplan. Oxford: Oxford University Press, 1975. Reprinted in [Cassam 1994]: 140-159. Page references to reprint.
- [Cassam 1994] Cassam, Quassim. *Self Knowledge*. Oxford: Oxford University Press, 1994.
- [Castañeda 1969] Castañeda, Hector-Neri. On the Phenomeno-logic of the I. *Proceedings of the 14th International Congress of Philosophy, iii*. Vienna: University of Vienna, 1969. Reprinted in [Cassam 1994]: 160-166. Page references to reprint.

- [Dennett, 1981] Dennett, Dan. Where Am I?. Chapter 17 of *Brainstorms*. Cambridge: MIT, 1981: 310-23
- [Descartes 1641/1984] Descartes, Rene. *Meditationes de prima philosophia*. Translated as *Meditations on First Philosophy* in J. Cottingham, R. Stoothoff, and D. Murdoch, trans., *The Philosophical Writings of Descartes*. Cambridge: Cambridge University Press, 1984.
- [Evans 1982/1994] Evans, Gareth. Self-Identification. From *The Varieties of Reference*. Oxford: Oxford University Press, 1982: 205-33. Reprinted in [Cassam 1994]: 184-209. Page references to reprint.
- [Frege 1918/1967] Frege, Gottlob, 1967a. The Thought: A Logical Inquiry. In *Philosophical Logic*, ed. P. F. Strawson. Oxford: Oxford University Press, 1967: 17--38. This translation, by A.M. and Marcelle Quinton, originally *Klogische Untersuchung*, appeared in *Beiträge zur Philosophie des deutschen Idealismus I* (1918): 58--77. Page references are to the Strawson volume.
- [Grice, H.P. 1941/2008] Grice, H.P., 1941. Personal Identity. *Mind* 50 (1941). Reprinted in [Perry, 1975]: 73--95.
- [Kaplan 1989] Kaplan, David. "Demonstratives". In [Almog et. al. 1989].

- [Lichtenberg, 2000] Lichtenberg, Georg Cristoph. *The Waste Books*. Translated with an introduction by R.J. Hollingdale. New York: New York Review of Books, 2000.
- [McTaggart 1921-27] McTaggart, John McTaggart Ellis. *The Nature of Existence*. Cambridge: Cambridge University Press, 1921—27.
- [Perry 2000] Perry, John. *The Problem of the Essential Indexical*, Expanded edition. Stanford: CSLI Publications, 2000.
- [Perry, 2001] Perry, John. *Knowledge, Possibility and Consciousness*, Cambridge: MIT Press, 2001.
- [Perry, 2002] Perry, John. *Identity, Personal Identity and the Self*. Indianapolis: Hackett Publishing, 2002.
- [Perry, 2008] Perry, John (ed.), *Personal Identity*, 2nd edition. Berkeley and Los Angeles: University of California Press.
- [Perry, 1975/2008] Perry, John. Personal Identity, Memory, and the Problem of Circularity. In [Perry, 2008]: 135--155.
- [Perry, 2009a] Perry, John. Directing Intentions. In Joseph Almog and Paolo Leonardi, eds., *The Philosophy of David Kaplan*. New York: Oxford University Press, 2009.
- [Perry, 2009b] Perry, John. *Reference and Reflexivity*, 2nd edition. Stanford: CSLI Publications, 2009.

[Shoemaker1968/1994] Shoemaker, Sydney. Self-Reference and Self-Awareness. *Journal of Philosophy* (1968): 555-67. Reprinted in [Cassam 1994]: 80-93. Page references to reprint.

[Wittgenstein 1958] Wittgenstein, Ludwig. *The Blue and Brown Books*. Oxford: Oxford University Press, 1958.