

THE SENSE OF IDENTITY

By John Perry

[Essay 11 of *Identity, Personal Identity and the Self*. Indianapolis: Hackett Publishing, 2002.]

1. THE PHILOSOPHICAL SELF

As I write this, I see a specific hand guides a specific pen across a specific page at a specific time and place. The hand belongs to John Perry---JP for short---one among the billions of persons who exist. I have a rather special relationship to JP, one which I can express by saying "I am JP." He is the one and only one among all the persons who ever have existed or will exist, who happens to be me. It is natural to take the special relationship to be identity; there is just one thing, one entity, one metaphysical unit, that is both the person I call "me" and the person I am calling "JP." We are the same not only in this possible world but in every possible world that one could describe or imagine, for there is only one thing to imaginatively project into different circumstances. So it seems that I am necessarily JP, and could be no other person.

But this doesn't seem quite right. In certain moments, there seems to be at least a trace of contingency in this relationship between JP and me. In fact, I see the world from his perspective, and interact with the world through his body. But I can imagine being someone else; having the perspective of another person, with a different body, living in a different place, perhaps even during a different period of time. I can imagine being Napoleon, exiled on Elba, or Bertrand Russell, puzzling about "the," or even Tom Nagel, writing the article from which these thoughts are drawn (Nagel, 1983).

These thoughts, that I am not John Perry, that I am Napoleon, or Bertrand Russell, or Tom Nagel, are surely false. But they do not seem to be utterly

incoherent or even necessarily false. They seem like possibilities. Indeed, they seem to be possibilities of a sort we need to think about. It is important for various projects, in particular literary and philosophical projects, to imagine being other people; people with different needs, emotions, political views, social roles, incomes, interests and so forth. What we are doing does not seem to be incoherent. We can first view the world from no point of view, and then imagine what it would be like to have different points of view than the one that we ended up with, almost, as it seems, by accident. We can try to think things through from the new, imagined perspective. It is not like trying to reason from the premise that $2 + 2 = 5$, or that there are round squares, or some other contradictory starting point. There are limits to what one can figure out by imagining that one is someone else, but they are more like the limits of reasoning from any other counterfactual premise, than the roadblock of contradiction.

Ordinarily, I think of the world from my point of view, centered on the here and now. We might call this our "subjective" point of view, but I prefer the term "agent-relative," as being somewhat less encumbered with various alternative and not quite on-target meanings. The agent is the thinker, the speaker, the doer. An agent-relative role is a role things play in the life of agents; different things for the same agent at different times and for different agents at the same time. What is in front of me is not what is front of you; what is to my left now is not what will be there tomorrow; and so forth. When we classify events, places, things and people by the roles they play in our lives, we are adopting the agent-relative perspective. Events are divided into past, present, and future. Places are spread out in various distances, those that are here, those that are there, those that are near, those that are far, those to the left, those to the right, those above, those below, and so on. As to people, there is me, the person I am, you, the person I am addressing, my family, my neighbors, my generation, and so on.

But I can also think of the world from no point of view, impersonally or objectively. Events are classified by their relations with each other: before, after, or simultaneous with. There is no "now" in this way of thinking; or rather there are many; each class of simultaneous events provides a "now," but none of these nows is by nature more privileged than any others. Places are organized by their relation to one another, or some absolute coordinate system. None of them is a "here" objectively; each is "here" relative to itself, "somewhere else" relative to other places. People are thought of by their names, perhaps, or their position in space and time. The representation, so long as it remains objective, has no "I"; I am in it, but everyone else who is in it is equally an "I," relative to themselves, and I am "someone else," relative to them. I realize that in this objective world, I am nothing special. The evidence of my senses counts for no more than anyone else's in arriving at truth; my interests and desires, pains and pleasures count for no more than anyone else's in considering what is right.

If we think that facts are objective, and the world is some sort of totality of objective facts, then it seems this view from nowhere is a more faithful representation of reality than any agent-relative view. And surely there is nothing it needs to leave out. If all the facts are objective, each can be represented impersonally. It doesn't leave out me; John Perry is in this view of the world, with all of his thoughts and experiences, desires and interests, pains and pleasures. But what of the fact *I am John Perry*? The fact whose contingency seemed so important? Where is it in the world? How is that fact to be represented? Can I find it in the representation? It seems it cannot be, for then the impersonal representation it would build in a point of view; it would not be a view from nowhere; it would represent my view, accurate relative to me, but not relative to you, for the thought "I am John Perry" is true when I think it, but false when you do. But...*I am John Perry*. This is some sort of fact; if not an objective fact, then what kind of fact is it?

The last thought, the thought that I am one of those people in the objective world, that I am JP, is difficult, perhaps impossible, to express without the first person. It seems (a bit) contingent that I am JP. It does not seem contingent that JP is JP. That JP is JP is clear in the objective representation, by the very fact that JP is there. But the fact that I am JP isn't in the representation; as we saw, if it were, the representation would cease to be objective. There is a problem here, which I'll call the problem of the philosophical self.

Contemporary semantics ought to help us to understand what I think when I think "I am JP." It would thus provide the sense of identity, and illuminate the philosophical self. But it is not clear how it does so. The standard semantics for indexicals and names identifies the propositions expressed by "I am JP" with the necessarily true proposition expressed by "JP is JP"---a thought anyone can think truly (Kaplan, 1983). It identifies the proposition I express with "I am Napoleon," with the necessarily false proposition that JP is Napoleon. But that is not the proposition I am thinking when I imagine being Napoleon. Something seems to have been left out, after all. But from what has it been left out? Is there a side of the world left out, by thinking that all facts are objective? Or is a part of ourselves left out, when we take ourselves to be just flesh and blood persons with a perspective on the world? Or could it be something more boring ---- something left out of the semantics we have for understanding and describing our thought and language?

2. *THE OBJECTIVE SELF*

One approach to this question is to find something special, other than JP, for "I" to refer to, at least my philosophical moments; something that is, or might reasonably be thought to be, only contingently related to JP. Thomas Nagel, whose thoughts we followed in setting up the problem, advocates this approach in his essay, "The Objective Self" (Nagel, 1983). For each person who has a conception of an objective or centerless world, there is an objective self. This self perceives the world through the person's sensory systems and affects the world

through the person's motor systems. But an objective self should not be identified with the person to whom it has this contingent connection. When the word "I" occurs in the expression of a philosophical thought of the sort described above, it does not refer to the person who expresses the thought. It refers to that person's objective self. When I thought, in this philosophical mood, "I am John Perry" the "I" referred to the objective self, the "John Perry" to an occupant of the world whose perspective the objective self has.

So Nagel gives the word "I," as used to express the philosophical thought, a new reference; the objective self, rather than the person. He also gives it a new sense; that is, a condition that identifies the reference: "the subject of this impersonal conception."ⁱ In my impersonal conception of the world, there is a representation of John Perry, a rather full and robust one, since I know a lot about him, or at least have many opinions. But its status, within the impersonal representation, is on a par with my representations of everyone else. Usually I have another rather intrusive representation of myself, one I might retain even if I forgot who I was, based on the contemporary information I can pick up about myself through feelings and perception, and tied to the word "I." But in this philosophical mood this representation is attenuated; I bracket off most of this information; I focus on the impersonal conception. But I cannot fully sever the connection; while what is conceived may be objective and impersonal, the mental conception itself belongs to only one of those people represented in it; that person has a special way of attending to it, which allows him to think of it as *this*. Whatever else I may be, at the moment of objective philosophical thought, when all that can be has been bracketed and put aside, *I* am the subject of *this* impersonal conception.

I like Nagel's suggested sense for philosophical uses of "I," but not his suggestion of a new reference. I think of myself as saving Nagel's insight from his metaphysics.ⁱⁱ

3. *NAGEL'S PROBLEM*

Nagel's asks: "How can a particular person be me," and he immediately glosses this question as follows:

Given a complete description of the world, from no particular view, including all of the people in it, one of whom is Thomas Nagel, it seems on the one hand that something has been left out, something remains to be specified, namely, which of them I am. But on the other hand there seems no room in such an objectively described world for such a further fact...everything true of TN is already in it. (Nagel, 1983, pp. 211-212)

This further fact is also described as "The fact that one of these persons, TN, is the locus of my consciousness, the point of view from which I observe and act on the world." (Nagel, 1983, p. 213)

It will be worthwhile to pause for a moment to get more of a feel of this problem. Let's imagine that Tom Nagel is standing next to an ordinary, if somewhat large, physical representation of a very large part of the world. I don't think we need to suppose that this is a representation of the whole world, for that is clearly a requirement that might bring large problems with it. Just a representation of a lot of stuff, let's suppose. We can imagine a huge, illustrated *Who's Who*, with a photograph of Nagel, and a listing of important events from his life and his accomplishments, and the same for a billion or so other people.

This is an objective representation, in that it contains no indexicals or demonstratives, the linguistic expressions of the agent-relative way of classifying things. That is, the content of the representation doesn't depend on who wrote it, when, and where; we don't need to know this information to understand what is said. An objective representation has the same content, no matter who created it or uses it.

To appreciate the difference, consider telephone books, great objective repositories of useful information. Here is a typical entry from one:

Chung, Dae-Hyun.....312-4312

There is no indexicality, no context sensitivity. There might be more than one Dae-Hyun Chung in the world, but this is an issue of ambiguity or something like it (See Perry, 2001b), or perhaps even different words (See Kaplan, 1990) and not indexicality. Once the words and meanings are clear, there is no part of the entry that directs us to context, the way "I" and "you" and "here" do.

We could have a telephone book that used indexicals. I could make such a book and give it to Dae-Hyun, with entries like,

You.....312-4312

Me.....372-5191

Imagine that a philosophy department, as a sort of experiment, or perhaps as a protest against realism and objectivism, puts out non-objective phone book, with entries like:

Me.....41275

Her.....45682

That guy over there73298

It probably wouldn't work very well. It would put a large cognitive burden on the user. To use it, one would have to learn and keep in mind facts about the creator at the time of creation of the phone book. Who was he---the "me" of the phone book? Who was he referring to with "her"? Who was he looking at when he wrote "that guy over there? And so forth. Once one found out these answers, one would naturally *annotate* the phonebook with names. Perhaps mine would look like

Her (Debra).....45682

That guy over there (Bratman).....73298

This seems to suggest that it is objective representations that are truly useful.

But consider this old joke. Dae-Hyun is talking to a woman he has met at a party. "We've had a nice chat," he says, "I'd like to see you again. Can I have your phone number?" "It's in the book," she replies. "And what is your name?" Dae-Hyun goes on. "It's in the book, too," she replies.

Now suppose that Dae-Hyun has an ordinary phone book with him. You can even suppose he has memorized it. All the facts about people in the town and their names and phone numbers are in his possession. They are objectively represented; no indexicals, no demonstratives. This won't do him much good. It seems objective representations aren't always what we need, either. If the woman was making a joke, and not brushing Dae-Hyun off completely, she might agree to annotate his phone book, by writing "me" next to one of the numbers, or pointing to a name and saying, "that's me." Then Dae-Hyun would have the information he needs. It's beginning to look like there is something really useful about the interaction of the two kinds of representations, as happens with these annotations. We'll return to this point below.

Back to Nagel. Let's suppose that he is looking at the page from *Who's Who*, and he recognizes himself as he reads the entry. He points to the entry and says, "That's me" or "I am (that) Thomas Nagel," or, for short, "I am TN." Now he has said something that expressed his discovery. Perhaps he annotates the objective representation. But *what* did he discover? It seems like it is a fact, but which fact? None of the ones that are listed in the entry, it seems, for he was able to grasp them without making the crucial discovery.

The point is not that there is some philosophical difficulty about how recognition happens. We can imagine there is a sort of pattern-matching with attributes of the newly presented objects and objects about which one already has information. When there are enough important matches, the two are identified. The problem is rather with what recognition means; the sense of identity, and in particular, the sense of self, of identity with the person doing the

identifying. What possible worlds does this identification exclude? What fact about the world does it represent? What fact is it, the grasping of which constitutes recognition?

Let's try to supplement the entry to see if we can find the fact. Suppose Nagel was reading copy 5 of *Who's Who* in the Princeton Library on August 20, 1983 at 3 p.m. So, how about adding, to the Tom Nagel entry, "Read, August 20, 1983, 3 p.m., *Who's Who*, copy 5, Princeton Library. This is an objective fact in the requisite sense. Have we added the fact the grasping of which amounts to Nagel's recognition that TN was him?

It seems we have not. For again, Nagel could read that entry and not be sure. Then, if he figured it out, he would say, "Oh, I am TN." Maybe he even looks at his watch and at the spine of the volume where it says "copy 5." Still there seems to be a gap between his reading and believing this additional fact about TN, and the realization he expresses with "I am TN".

Let p be Nagel's current perception of the TN entry. Add to the entry, "Has perception p on August 20, 1983." So now Nagel learns from the entry that TN has p which is the very perception that he is having. How can he believe that, and not believe that he is TN? But of course he can, for he may not realize that the perception referred to by " p " is the very one he is having. The realization "This perception I am having is p " presents the same problem again. He could already know that p is the perception TN is having without having, without realizing this.

The pattern is clear; any fact, an objective representation of which we might add, seems not be to the fact in question. Nagel's grasping that fact, via the objective representation, will not constitute his realization that he is the TN whose representation he studies, for that fact could always be added to the objective representation, and grasped by Nagel, without his having that realization.

4. *AGAINST THE OBJECTIVE SELF*

Can the objective self rescue us? According to Nagel, for each person there is an objective self, which is contingently related to that person. So for TN there is an objective self, we can call it "OSTN." And for me there is one, we can call it "OSJP." These objective selves have no specific location in space and time, but they do have a special though contingent relationship to the body of the person whose objective self they are. When one has the philosophical thought, "I am TN" or "I am JP," the "I" has the sense "the subject of this objective representation," and stands for one's objective self.

It is very difficult to see how the postulation of objective selves provides any solution whatsoever to the original problem. Part of that problem was to find what to add to our objective representation of the world, to correspond to the fact that TN discovers, when he discovers that he is TN. Now we can add our objective selves to the representation, and it doesn't seem to help at all.

When TN looks at this representation, he gets the information that OSTN belongs to TN and, if he looks in the right place, that OSJP belongs to JP. He can know this, and still remain confused about whether he is JP or TN. When he realizes that he is TN, and thinks, "I am TN," he is realizing something new that he didn't realize already, even though he did realize that TN was the one with OSTN. So it is very difficult to see how postulation of objective selves helps at all with our problem. There is nothing in Nagel's explanation of objective selves to explain why TN's having an appropriate contingent relation to one objective self (OSTN) rather than another (OSJP, say), isn't an objective fact, that can be represented in an objective way.ⁱⁱⁱ But as we noted, any time a fact is a candidate for being what Nagel grasps, when he grasps that he is TN, it turns out not to work. We add an objective representation of the fact to the objective representation that Nagel is studying; he can in theory grasp the fact, without grasping that he is TN. This argument seems to disqualify the relation between

TN and OSTN from being what Nagel grasps, as surely as it disqualified all of the other candidates we have considered.

5. *THE SUBJECT OF THE IMPERSONAL CONCEPTION*

The new reference for "I" doesn't help. How about the new sense: "the subject of this impersonal conception of the world." Nagel says, having introduced objective selves:

I believe this accounts for the content of the philosophical thought we have been trying to track down. It is qua subject of this impersonal conception of the world that I refer to myself as 'I' in thinking the philosophical thought, "I am TN." Though the 'I' is still essentially indexical, the content of the thought is that this impersonal conception of the world is attached to the perspective of TN and is developed from that perspective...while it does not translate the thought into one about the world objectively conceived, it does identify an objective fact corresponding to the thought..."

The content for the missing thought seems to be:

The subject of this impersonal conception is TN

Let's acknowledge that when we have the philosophical thought, we are indeed thinking of ourselves in this way. So when I have my thought, I have an objective conception of the world. Now we are not thinking of a huge phone book or some other public representation, but my own internal objective, impersonal representation of the world. It is my conception; I am the subject of it. And so I think of myself in this way, and when I think "I am JP," I am thinking "the subject of this conception of the world is JP." Still, how does this solve the first problem?

It doesn't seem to solve it at all. My impersonal conception of the world can be added to the objective picture. We can add representations of a variety of impersonal conceptions; TN and JP are both subjects of impersonal conceptions,

and we can represent those conceptions and their relation to TN and JP in the picture. And either TN's or JP's conception or both might conform to our picture. I might be thinking, "So the world has TN in it, and he has an impersonal conception of the world, call it "ICTN." It has JP in it, and he has an impersonal conception of the world call it "ICJP." Now my use of "this impersonal conception" and "ICJP" both refer to the same impersonal conception, mine. And I can grasp that JP has ICJP, and still wonder: am I JP?

But we must remember that Nagel said that when I have the experience of the philosophical self, I am thinking "the subject of *this* impersonal conception is JP," not "the subject of ICJP is JP." Do these two thoughts correspond to the same fact or not? If they do, then the fact that the subject of ICJP is JP can no more be the one I grasp when I realize that I am JP, than any of the other facts we have considered.

Nagel said the thought was "essentially indexical," meaning that the "this" was crucial. Are there then somehow different facts corresponding to the two formulations, "the subject of ICJP is JP" and "the subject of *this* impersonal conception is JP" ?

There are. But to find them, and put Nagel's insight into a place where we can say why it works, we need to get less profound for a while.

6. INFORMATION GAMES

I want now to go back to those annotated representations that we found so useful when we were thinking about phone books a couple of sections back. Let's shift to a slightly more lofty example, business cards.

Suppose I am trying to raise money for Stanford and I am to meet a potential donor. First I come into the room and see a well-dressed person in front of me. Remembering my coaching from the Development Office, I smile, walk over and extend my hand and say, "Hello, I am John Perry."

What happened here? There was an object in front of me. I have a technique for finding out about objects in front of me: I open my eyes, look straight ahead and notice things about them. I used this technique and discovered that there was a nicely dressed important looking person in front of me.

I also have techniques for doing things to persons that are in front of me. I can introduce myself by looking at them while I say, "I am John Perry". I can offer to shake hands by moving towards them and extending my hand. These are appropriate if there is a person in front of me--- but not if there is an apple or a wall or a hungry bear in front of me. There are different things I know how to do with apples in front of me. I can take a bite out of them by executing a certain complex movement: extending my arm, grabbing them, moving them to my mouths, etc. The same movement probably would be a way of irritating a hungry bear.

Being in front of is what I called an *agent-relative role*; it is a role an object can play in someone's life. With it are associated *epistemic* and *pragmatic* methods---methods for finding out about the object and methods for doing things to the object (more accurately: doing things the success of which depends on facts about the object). I will speak of picking up information *via* a role and acting *via* a role.

When I introduced myself and extended my hand, I was applying information I picked up *via* the *being in front* role, to guide an action I performed via that very same role. This is a very basic *information game*. I'll call it the "straight-through" game.

But now let's return to my story: the potential donor gives me her business card. This card is full of objective representations: her name, Sarah Toorich let's say, her company, her position in it, and so forth. There is nothing on the card indicating her relation to me.

When she leaves, I put the card in my shirt pocket. For the rest of the day I carry around this objective representation---objective in the sense given above, there is nothing in the card that makes its content depend on who created it. It does not contain words like "I" and "here" and "you".

In the beginning, the information I had about Sarah Toorich, the way she looked and where she was in relation to me, was drawn from my perception of her as the person in front of me. When she gave me her business card, I had objectively represented information but it remained *attached* to that perception. Attached information is like what I called annotated information above, except that there is no bit of language serving as an annotation. Instead of having the annotation "that person over there" attached to the business card, I simply have it connected with my perception of Sarah.

When I left the room, I retained the objective representation, the business card, but it was detached from any perception. It was no longer tied to a perception of Sarah Toorich. What possible use can such detached information -- - objectively represented information --- be to me? It can be useful to me if I am later in a position to interact with Sarah and can recognize her. Once I have done that, I reattach the objective representation to a perception, and use it to guide my action.

Later on I run into this same person and I say, glancing at the card, "Hello Ms. Toorich. How are things going at the old Megbux.com today?" In this case I apply information I got in one situation, via one role, in a new situation. I will call this the *detach and recognize* information game. "Detach" because I detached the information about Ms. Toorich from any agent-relative role. "Recognize," because in order to apply the information, to use it to guide my action, I had to encounter her again and to recognize her, in order to re-attach and use that information already "on file".

The detach and recognize information game is very much a part of human life. It works because we live in a world with certain very general

characteristics---a "Strawsonian" world, we might say, remembering some of the themes of Peter Strawson's *Individuals* (Strawson, 1959). The success of our actions often depends on the relatively enduring properties of objects. We are often in a position to gather information about objects that will be relevant to the success of actions we may want to take when we encounter them in other situations. And we are often able to recognize objects we encounter more than once. Put these characteristics together, and we have a world in which the detach and recognize information game is a very useful technique.

The detach and recognize information game also fits into a more complex information game, involving communication. What we want to communicate in many cases is detached information, that each party to the communication game can use to recognize and attach to the particular perceptions they have of the object the information is about.

This detach and recognize picture should help us put objective representations in their proper place. They are basically *supplements* to agent-relative ways of thinking and acting. They give us information that helps us get ourselves into position for such thought and action. Some of the information that helps us recognize who we are interacting with : my memory of what Sarah Toorich looks like, the aids to recognition I jotted on the back of the card. Other information that helps us decide what to do once we are interacting: she is a wealthy but pensive business woman, so ask for a contribution to the Accounting and Philosophy joint major.

Although the world is objective, all perception is by agents at times and all actions are performed by agents at times. An objective representation with no possibility of being reattached to its source via some agent-relative role that supports pragmatic techniques is quite useless. If I can't recognize Sarah Toorich, for example, the bits of information about her I have stored in my shirt-pocket, or my palm pilot, or in the little mental 3 x 5 note-card that I like to think comprises my internal objective conception of her, will be of no practical value.

Philosophers often think of beliefs as relations to propositions, which are objectively true or false, not true from one perspective, false from another. In describing the belief, we may identify the constituent of the proposition indexically, but that doesn't mean we think that those indexicals play a role in his belief. If I say, walking in Manhattan, "Nagel believes this city is exciting" I am saying what Nagel believes, but not how he is thinking of it. He might be in London as I speak, unable to think of New York City as "this city." But what I say is true, for he has an objective way of thinking of New York City, and associates the property of being exciting with it. This part of the picture of belief, as relations to is at the heart of the ordinary philosophical way of thinking of belief as a "propositional attitude."

This picture gets at only part of the story---the detached, objective, impersonal part. It can't handle the derivation of and re-application of detached beliefs. The problem that Nagel finds, that leads to the postulation of the objective self, is just the problem of this missing element in the philosophical picture of belief merely as a propositional attitude. This is the point I'll now try to make vivid and plausible, going back to our exciting example.

Suppose I am in a meeting with Ms. Toorich and four or five other business persons around a big table. I have been very organized and I have all of their business cards in front of me. Unfortunately I don't know who is who. I jotted down some "aids to recognition" on the back of the cards when I received them, but I have momentarily forgotten that I did so. Unless I can *orient* myself the information is useless. That is, I need to reattach my objective information, my business cards, to the agent relative roles that are occupied. That is my theory-laden way of saying I need to recognize the people in front of me.

Now suppose that the person across from me is Sarah Toorich. Echoing Nagel, in what objective fact does this consist? When I recognize *her* as Ms. Toorich, what fact do I grasp?

This question has a false assumption, namely, that recognition *ever* consists in grasping an objective fact, or rather consists in that *and nothing more*. This is simply not so. Recognition is quite a different concept than belief in a proposition. Recognition consists of attaching objectively represented information to some perception, readying the information for use. Recognition will occur *because of* the beliefs one has; one compares the attributes of the objects one perceived with the attributes of the objects one has on file. But recognition does not *consist in* having those beliefs. I recognize Sarah Toorich as the person in front of me, or as the person on the phone, or as the person I see in the distance. Recognition is a prelude to action that is then partly guided by belief, a prelude to the application of detached information (or misinformation, for that matter).

7. *SELF-RECOGNITION*

I too have business cards, that say on them that my name is "John Perry" and that I am a Professor of Philosophy at Stanford and give my phone number, email, and etc.

Suppose now that on my way to work one morning I fall off my bike and hit my head and get temporary amnesia. I'm sitting on the curb, wondering who I am. I reach in my pocket a pull out the business cards I have there, searching for a clue. Here is a card that says, "Sarah Toorich, Megabux.com." Could that be me?

I think as a matter of fact people with amnesia don't forget whether they are male or female. I'm not sure. Suppose that in this case I have. I can check to see if I am a male, say by feeling, say, the beard on my face. This is an agent-relative way of finding this out, quite different than how I would check to see if you were a male, which I could usually do just by looking. I am wondering if I am Sarah Toorich. I eliminate this possibility since "Sarah" is usually a woman's name.

Next I pull out one of my own cards, "John Perry". That is a bit more plausible---but who knows?

I want to analyze this situation much like I did the situation where I was sitting with the business people around a table. I had plenty of objectively represented information about the people in various positions relative to me, but wasn't able to apply it, since I didn't recognize them. That is, I couldn't attach the detached representations I had to the agent-relative roles they occupied, and thereby make some use of the information on the cards.

I claim that identity is an epistemic/pragmatic relation self is an agent-relative role, consisting of identity with the agent. This is elaborated in Essay 10, so I'll be brief. Identity is a relation that brings with it certain epistemic and pragmatic techniques. There are special ways of getting information about the person one has the relation of identity to---one's self--- and special ways of acting so as to have effects on that person. We are happy when children learn to recognize when they need to go the bathroom. We are happy when they learn to feed themselves. It's quite a different ability than feeding other people, which we usually discourage children from doing too much of.

To get back to me. We left me sitting on the curb after a bike accident, wondering who I am, staring at my own business card. That's where I am in the example. In the argument, I am comparing this situation with that earlier one, where I had objectively represented information about several people around a table, but couldn't apply it.

I have just finished making my first point about that, which is that *self* is an agent-relative role, being identical with the agent, that is associated with special epistemic and pragmatic methods---just as *being in front*, *being on the phone* and other agent relative roles are.

My second point is that in our own case, as in the case of others, objective representations of facts are useless in and of themselves, and don't ever explain any action. Their importance is as potential parts of attached representations.

This is just the situation I am in, as I sit on the curb. My business card gives information about me. Lots of people have this card. For each of us, it is of interest as a potential component of applied information. Ms. Toorich has my card. Next time she sees me, if she recognizes me, she can apply this information to me, greet me with my name, make some wisecracks about the philosophy business, and so forth.

I have here in my hand this same objective information, this same potentially attachable information. It would be very useful for me to know my name, etc. But before I can use the information, I need to recognize myself, to associate the information with my self-notion. Luckily, being a man of very orderly habits, I jotted down a few salient facts about the appearance of a person on the back of these business cards to help me recognize them, and I even did this in my own case. So I know that John Perry has gray hair and a gray beard and wears glasses. I use self-informative techniques for getting information in this case, just as I use persons-around-the-table techniques in the other. For example, I can see a little bit of my moustache, just by looking down towards the floor, and I can see that it is gray. I can also see the rims of my glasses. So I figure out who I am, call home, find out what kind of health insurance I have, and check into the hospital until my amnesia clears up. When I check in, I will consult my business card and say, "I am John Perry. I teach philosophy. My phone number is 327-9161." These remarks will be guided by the card, because I have attached it to my self-notion, my repository for information acquired in normally self-informative ways, and other information, gotten in more neutral ways, like business cards, that I have determined to be about me (See Essay 10).

8. *THE MISSING FACTS*

Let's review the bidding. We've got the philosophical uses of "I," and we want to know what their content is; what I grasp when I think, in the right, Nagel-inspired mood, "I am JP." The strategy is to see this as a special case of self-recognition, and self-recognition a special case of recognition, and recognition as

involving an interplay of the objective and the agent-relative. I've developed a certain picture of recognition, and applied it back as far as the ordinary cases of self-recognition. But I haven't really said what the missing facts were.

As I sit on the curb, looking at my business card, I am in a situation in many ways analogous to Nagel's problem situation. I'm looking at an objective representation of me, my business card with my name and some essential information about me, all objectively represented. As I stare at it, I learn all of this stuff about JP. He is a philosophy teacher at Stanford, has a nifty email address, and the like. So I know all of those facts about JP. Then I realize I am JP. What does this knowledge consist in?

Here is where modern semantics seems a bit unhelpful.^{iv} Since the 1970's we have been told that names and indexicals are "directly referential" or "rigid designators." That means that "I," as said or thought by me, and "JP," as a name that stands for me, contribute exactly the same thing to the propositions in which they occur, namely me. Suppose you say "John Perry was born in Nebraska." What does the world have to be like, for what you say to be true? Do I have to be named "John Perry?." It seems not. After all, it was true that I was born in Nebraska before I was named John Perry. I would have been born in Nebraska, even if I had been named "Elwood Perry" or something like that. Although you exploit the fact that I was in fact named "John Perry" and not "Elwood Perry" when you say "John Perry was born in Nebraska," it is not part of what you say. What you say seems to be just that a certain person, that in fact is named "John Perry," was born in Nebraska.

Now suppose I say "I was born in Nebraska." What does the world have to be like, for what I say to be true? Well, it would have been true, even if I hadn't said it. My saying it didn't have any effect on where I was born. The word "I" stands for me because I used it, but that isn't part of what I said. What I said would have been true if I hadn't said anything; what I said was true because a certain person was born in Nebraska, period. But then what I said

with my utterance of "I was born in Nebraska" was just what you said with your utterance of "John Perry was born in Nebraska." The truth-conditions are just that a certain person, me, was born in Nebraska, the same for both statements. That's more or less what philosophers mean by saying that both names and indexicals contribute the thing they stand for (in this case, me) to the propositions they express.

But then

- 1) JP is JP
- 2) I am I
- 3) I am JP

all seem to say the same thing, to express exactly the same proposition, the trivial and necessary proposition that John Perry is John Perry. But only 3) expresses recognition.

Here is where I think modern semantics needs some supplementing, which is relevant to our problem. The basic point is a distinction, or rather a pair of them, one for language, one for thought:

What is said by a statement, versus the truth-conditions of the statement;

What is believed in virtue of having a belief, versus the truth-conditions of the belief.

Let's go back to your utterance of "John Perry was born in Nebraska." We agreed that it is not part of what you said that I, or anyone, is named "John Perry." You are talking about people and states, not names. Still, your actual utterance would not have been true, if I hadn't had that name. Actually an enormous number of things have to be the case for your utterance to be true. We just don't count most of them as part of what is said. What we are *usually* interested in is the conditions put on the subject matter---the things the names

and indexicals stand for. That's what we count as "what is said." But the other truth-conditions are very relevant to understanding how communication works.

Suppose, for example, that a drunk in a bar says something scurrilous about Nebraska, and then asks in a loud voice, "Is anybody here from Nebraska?" Trying to be helpful, you say "John Perry is from Nebraska." The drunk will learn that someone from Nebraska is named "John Perry." That wasn't what you *said*, but it was something that had to be true for your statement to be true---not for *what you said* to be true, but for the *statement you actually made*, your utterance, to be true.

Still things aren't too bad. He knows that someone from Nebraska is named "John Perry," but he doesn't know that I am. On the other hand, if I, trying to be helpful, had said, "I am from Nebraska," even though I would have said the same thing, things would have been a lot worse. He would have learned that someone right in front of him, that he could reach out and hit, was from Nebraska. That wouldn't have been part of what I said. But he knows that for my utterance to be true, the person "I" refers to must be from Nebraska, and he knows that "I" refers to the person who uses it, and he saw that I used it. I would have been in trouble.

Here we have two ways of saying the same thing, and yet they have different *total* content. The facts that explain the difference are part of the truth-conditions, but not part of the *subject-matter* truth conditions. They are conditions on the things we take for granted in getting to the subject matter, namely, the words *themselves*. For this reason I call them *reflexive* truth-conditions or *reflexive* content.

This same distinction carries over to thoughts. In Essay 10 I discussed self-notions. Notions are ideas we have of things, and self-notions are ones that are tied to the epistemic and pragmatic methods tied to identity. The self-notion is the repository of information picked up in self-informative ways, and the motivator of self-effecting actions. I said that my self-knowledge involves my

self-notion, and is to be distinguished from mere knowledge of the person I happen to be. When I was sitting on the curb reading John Perry's business card, I had knowledge about John Perry, the person I happened to be. Self-recognition consisted in linking that idea of John Perry with my self-notion; I came to believe not only that John Perry worked at Stanford, but that I did. Before the episode of self-recognition I believed the proposition that John Perry worked at Stanford, and this is what I believed after the episode. But I believed it in a different way. Call my self-notion $self_{jp}$. My later belief can be true only if $self_{jp}$ belongs to someone who works at Stanford. That is the reflexive content of my belief, and it is this that changed when I recognized who I was.

Now let's turn to more or less normal cases of self-recognition. Suppose I am watching a home movie made in the late forties by my uncle Art. It is a movie of some kind of family reunion, and there are a lot of kids in the movie---my cousins. At one point I recognize that a certain small child is myself. "That's me," I exclaim to the generally uninterested audience.

Let's suppose the child in the movie is wearing a cowboy costume and playing with guns. Suppose further that I don't remember playing with guns, and didn't even realize that I ever had played with guns. Because I recognized the five-year old playing with guns as me, I added to the information associated with my self-notion that I used to play with guns. This addition is what motivates my saying "I used to play with toy guns!" Note that that the word "I" gives us a self-directed way of asserting things; the sentence I used is a way for any speaker to assert of himself, that he used to play with toy guns.

In Essay 10 I discussed knowledge that is acquired in normally self-informative ways, like one's current state of hunger, or the nature of the scene before one. It is the job of the self-notion to handle such things. As a result of recognizing that I am the five-year old in the movie, I include a bit of information gotten in a way that is not normally self-informative with information that is.

The person that I know in a special way to be hungry and tired, I also know to have once played cowboys and shot toy guns wildly at his cousins.

To get at what I learned, we again need to appeal to the reflexive level of the content of my thoughts. Before the recognition occurred, I had two beliefs which we can describe structurally:

The idea of being tired-and-hungry was associated with n_{SELF} .

The idea of playing with guns was associated with $n_{\text{kid in movie}}$

The subject matter contents of these beliefs imply that a single person was tired, hungry, and played with guns. This is because the two notions are of the same person, me. But this identity is not reflected in the structure of my beliefs, and so cannot affect my behavior. In particular, I won't say "I played with guns as a child." Once recognition has occurred, the information is associated with my self-notion,

The idea of playing with guns is associated with n_{SELF}

and I will say that I played with guns as a child, if asked.

What then is the sense of identity, of self, in this kind of situation? It is the complex of epistemic and pragmatic relationships that are most closely and firmly tied to the self-buffer. "The child playing with guns is me" is true if the person whose present perspective is *this one*, whose present sensations are *these*, whose present thoughts are *these*, and the like.

The core of our self-concepts, our sense of identity, our sense for "I," is as the knower of facts about objects that are playing agent-relative roles with respect to us, and the agent of actions that are done agent-relative ways. I am the possessor and controller of these hands; the subject of these sensations; the maker of these movements; the sufferer of these pains; and so forth. Since only we can attend to our own inner sensations and thoughts, and only we can see

our bodies and things around us from our perspective, it is natural to use the demonstrative "this" to express the aspects of our self-concept.

Russell once held a "hidden description" theory of the self, with similarities to the view being put forward here (Russell, 1956). "I" was a "hidden description" for "the person with *this* sensation," where the "this" signifies an act of internal attention. (Russell, Russell also held a "hidden description" theory of proper names; a sophisticated version of this genre of theory was developed by Searle, a "cluster theory." No *one* description corresponds to a proper name, but a cluster of weighted descriptions (Searle, 1958) .

My view of our self-concepts is something like a cluster version of Russell's hidden description theory of the self, in that I think we have a cluster of things in our self-concept, which are weighted in their importance to us. The most important and inseparable from us are the things in our own mental life that we can attend to and think of with an internally directed "this." In spite of this similarity with hidden description and cluster theories, however, my view is quite different on the crucial matters of reference and truth. My self-notion is a notion *of* me because it is *my* self-notion; that is, (a) it is a self-notion, one whose informational role is as the repository of information gotten in normally self-informative ways, and that motivates normally self-directed actions; (b) it is mine. It is *of* me even if it is full of false stuff. My self-concept or self-file, the notion together with the ideas associated with it, is *of* me because the notion is of me, not because I am uniquely denoted, or denoted at all, by the combination of ideas.

It seems to me quite likely that I am in fact denoted by some weighted combination of the properties that seem to me to be most certainly mine. Such properties are the core of our self-concept, because these properties are very hard to separate from ourselves. The degree of difficulty is not uniform, however. We can easily imagine cases in which certain hands are not mine, even though I see them in the peculiar way that people normally see their own hands. We can

even imagine cases in which the stomach the state of which I learn about through my feelings of hunger is not mine. These cases are not likely to occur, and perhaps only a philosopher would have the curiosity and patience to think (however superficially and incompletely) about the rewiring of nerves and the like that could produce these results. These are open only to philosophical doubt.

The separability of my own sensations and thoughts is at least an order of magnitude more difficult. Can we imagine any circumstances in which the sensations I can attend to are not mine? Where the thoughts that run through my head are not mine? John McTaggart criticized Russell on the grounds that the hidden description "this sensation" may be "improper," like "the senator from Nebraska," which doesn't denote anyone, since there are two senators from that state. I may be certain that I am one of the people with *this* sensation, but how do I know that I am the *only* person with this sensation, McTaggart asks? (McTaggart, 1927, vol. II, p. 67) On Russell's own theory of descriptions, "the person with this sensation is John McTaggart" would be false, if more than one person had the sensation in question. How can he know for sure, McTaggart asked, that no one else has the sensations he can attend to?

This isn't a problem for my account for two reasons. First, unlike Russell, I'm not trying to build up a hierarchical account of our knowledge with self-knowledge somewhere near the bottom. I'm as sure that I am the only person having my sensations as I am that I am the only person with my body and brain. Moreover, even if someone else could also have my sensations and thoughts, it wouldn't be a problem, since I don't need to salvage a proper description from my self-concept to have it be a self-concept of *me*. My self-beliefs are a cluster; but it is not the denotative properties of this cluster that make the beliefs about me. So even if against all odds I am not the unique person with these sensations and thoughts, or even the best fit for the weighted

sum of the most central parts of my self-concept, it would still be a concept of me.

Suppose, for example, that David Chalmers is right about qualia, and so I am wrong. (see Chalmers, 1996, Perry, 2001a). When I think "I am a philosopher who is right about qualia," I think something false; the indefinite description fits Chalmers, not me, even though this is a very important part of my self-concept. Suppose further that , as a result of his researches on the mind, he has learned how to share his thoughts and sensations with others when he wants, so that his thoughts are accessible to them in just the way their own thoughts are, and seem like their own thoughts. He often picks on me, transmitting thoughts via my mouthful of gold fillings, confusing and embarrassing me in various ways with the thoughts and sensations of a younger man. Perhaps on Thursday I am thinking about myself, and the central-most part of my self-concept is as a person who is right about qualia, and has *these* thoughts and *these* sensations. But in fact this is a moment when most of the sensations and thoughts are Chalmers'; they are either not mine at all, or not mine uniquely, and at least are more his than mine, since the sensations are caused by his bodily states in the normal way, and he controls the thoughts. So just considering the sensations and thoughts, my self-concept fits him better than it fits me. When we add that he is a philosopher who is right about qualia, an important part of my self-concept, he is clearly a better fit. Still, my self-concept is a very confused concept of me, not an accurate one of him.

This is of course in several ways a very fanciful example, certainly implausible, perhaps utterly incoherent. I do not claim to have a clear idea of the difference between Chalmers's causing me to share *his* sensations, and Chalmers's causing me to have sensations and thoughts that are mine, but are like his; I suspect there is no difference in what is conceived, and the latter is the correct way to describe it. But I'm not completely certain. Perhaps some day I'll have a clear conception of the difference, or a clear idea of why the first option makes no

sense. The bottom line is that even if it does make sense the way I described it, and my self-concept fits Chalmers better than it does me, it is still a concept of me, and I am thinking about myself when I use it. On my view there is need be no metaphysical necessity connecting me with some aspects of my self-concept, nor do I even have to be the one it fits best, for my self-concept to be of me. If there are such metaphysical necessities, and a good fit, that's fine. But it is not what makes my self-concept of me.

9. *CONTENT AND CAUSE*

My account has two sides, and it is not yet clear how they fit together. On the one hand I have said that objective knowledge is incomplete; it is made to supplement agent-relative knowledge; Nagel's problem is trying to find an objective proposition, belief in which would give him the requisite self-knowledge, but this cannot be done. On the other hand, I seem to have offered a proposition that does exactly that. I said that the difference between the thought expressed by, say,

(4) JP teaches philosophy

and that expressed by

(5) I teach philosophy

is that the belief expressed by (5) has the truth condition,

(6) The owner of $n_{SELF-JP}$ teaches philosophy,

where $n_{SELF-JP}$ is my self-notion. But (6) is an impersonal representation of an objective fact that could be added to the impersonal representation. What is going on?

The point that must be kept in mind is that (6) is not an additional belief that I gain. Most people don't have any *beliefs* about their self-notions. They don't have the requisite concepts to have these beliefs. One doesn't need to have

beliefs about one's self-notion, in order to have beliefs whose truth imposes requirement on it.

Another fanciful example will perhaps make this clear. Suppose I am giving a lecture about this theory to a group of students who accept everything I say as very plausible, if not absolute gospel. I give my own self-notion a name, " $n_{SELF-JP}$," and write on the

(7) The owner of the self-notion $n_{SELF-JP}$ is from Nebraska.

I believe (7); I have the requisite concepts to have such an odd belief. But believing (7) is not the same as believing that *I* am a philosopher. Each of the credulous students in the class will believe (7), but they won't thereby believe that they are from Nebraska.

The picture is this. There is a certain way of believing things, that involves the self-notion. It is a species of attached beliefs, beliefs that involve the kinds of notions I call buffers, that are tied to epistemic/pragmatic relations. What is special about these notions, and the beliefs that contain them, is the way they work, the way they are connected to our perceptual and motor systems. So the beliefs that involve these notions will have subject-matter content, but also they will have a different information handling role, and hence a different causal role, than other beliefs with the same content.

But their content and their causal role need to mesh. The subject matter of a belief involving a self-notion will be the person that owns the notion; that's who such beliefs are about. It will be true if that person has the requisite property. I have a belief, that involves the association of the idea of being from Nebraska with my self-notion. The subject matter truth conditions of this belief is that John Perry is from Nebraska, because John Perry is the owner of the belief. That content does not guarantee that it is a self-belief, however.

What guarantees this is the self-notion. Because it involves a self-notion, its truth puts conditions on that notion. The truth of a belief always puts

conditions on the ideas that make it up, just as the truth of an assertion puts conditions on the words that make it up. These conditions are not *what is believed* or *what is asserted*. But they are features of a belief's truth-conditions that tie into its structure. The level of reflexive content is where content and causal role meet.

Now consider (7). Belief in (7) will never explain much of anything. But having a belief with (7) as its reflexive content will explain a lot. Anyone can believe (7), but only I can have a belief with (7) as its reflexive content, because only I can have a belief with my notions as a components. The fact that being from Nebraska is associated with my self-notion explains why I say, "I am from Nebraska," when asked where I am from. That sentence provides a self-directed way of saying that someone is from Nebraska. If someone opens a booth given \$100 to Nebraska natives (not a very likely occurrence) I would rush over and claim my money. Each of the Nebraskans lined up at the booth will have a belief with a reflexive content like (7), but involving their self-notion instead of mine. In each case, it's not what they believe that explains their doing this, but how they believe it: via their self-notion. But if they believe that they are from Nebraska, via their self-notion, their belief cannot be true unless their self-notion belongs to someone from Nebraska.

Let's return to Nagel and the giant illustrated *Who'Who*. We add to the TN entry,

(8) TN is the owner of $n_{SELF-TOM-NAGEL}$

This obviously wouldn't help. Nagel could believe (8) without believing that *he* was TN. What he cannot do is have a belief with (8) as its reflexive content, without realizing that he is TN.

On this theory, then, there is no need for any kind of facts other than objective facts, no need for true propositions that are "subjective" or agent-relative. The facts that make it the case that Tom Nagel realizes he in TN are all

objective facts. But the propositions that need to be true for him to realize this, are not all ones he needs to believe, and the propositions he believes when he realizes this, need not include all of those that need to be true for him to do so. The reflexive content of his belief needs to be true, but he need not believe it.

As far as I can see, all facts are objective; I'm not very clear what it would be for a fact not to be objective. The proper place for the term "objective" to draw a contrast seems to be in the context of belief, knowledge, and representation. I do not think that all knowledge is objective, or impersonal, or should it be, or could it be. If knowledge is to be useful, the objective knowledge must be supplemented by agent-relative knowledge.

10. THE OBJECTIVE SELF

Nagel in his philosophical mood attends to a certain feature of his inner life, an impersonal conception of the world, which includes an impersonal representation of TN, with as many impersonal representations of facts about him as one may want. The inner conception labelled "TN" is a notion of Nagel, but not Nagel's self-notion. That is, Nagel does not apply the information associated with TN to his own situation, and he does not add information gotten via self-informative ways of knowing to the concept of TN. What makes the TN notion of a notion of Nagel is not that it plays the self-notion role in Nagel's life, but that it is in fact used as a repository for information about Nagel, gotten in other ways. In this Nagel's TN notion is like his notions of other people in the world. His knowledge of TN via his objective representation, before the moment of recognition that he is TN, is what we called "knowledge merely of the person he happens to be" in Essay 10.

At this moment, Nagel's situation is formally analogous to mine, sitting on the street corner after my bike accident. I have a sheaf of business cards. It seems to make perfectly good sense to hope that I am this fellow, hope I am not this other fellow, wish I were this third fellow, although I'm sure I'm not. Similarly, it seems that Nagel, because of philosophy, has severed the connection

between his self-buffer and his Nagel notion. It makes sense for him to imagine being various people, to be glad or unhappy he turns out who he turns out to be.

In this mood, one wants to say that there are many people I might be, that I might have been, that I might turn out to be. There is a sense of contingency in Nagel's being Nagel, and me being John Perry. The saying, "There but for the grace of God go I," can just mean that John Perry could have had a lot of bad luck, and wound up as a pauper or a tax lawyer, rather than having the privilege of earning a living as a philosopher. But it might also have a different meaning, that John Perry might be just as fortunate as he is, but *I* might not be; I might be someone else, doing something else, even tax law.

For familiar reasons, these thoughts are hard to make sense of. Since I am John Perry, there is just one person, one thing, one metaphysical being, that is both John Perry and I. When I take John Perry and myself to another possible world, I take only one thing. So I can't very well manage to find a possible world in which I am there, and he is not, or I am there, and he is there, but I am not he. Identity is a necessary relation; if *a* and *b* are one in any world, they are not two in any world. So what possibilities can correspond to these various thoughts we have been considering; what can we make of the felt sense of contingency we have in these philosophical moods?

11. *SEARCHING FOR CONTINGENCY*

For a contingent connection, we need two things. Here we have Tom Nagel attending to his own impersonal conception of himself, which we label "TN." There is not a contingent connection between Tom Nagel and TN, because there is just one thing that is both Tom Nagel and TN. There is no room for subject matter contingency.

How about the issue of how the thought fits into the world? Can we find that kind of contingency? That gives us two more things to deal with, the thought "I," which has the sense, in this situation, of "the subject of *this* objective

conception," since everything else is Nagel's self-buffer has been bracketed off. This is an agent-relative role, that Nagel in fact plays; he is the owner of both the objective conception, and of the thought directed at the objective conception, which is the sense expressed by the word "I" in this case. Can we find a world in which we have that very thought, that in fact is Nagel's, and Nagel, but Nagel is not the owner of the thought, and hence not the subject of the conception attended to? I don't think so.

That leaves us with one thing to work with, the link between Nagel's impersonal conception of himself, and Nagel. It is a conception of Nagel. But does it have to be a conception of Nagel? Is there a possible situation in which this very conception, the very one Nagel attends to, is *not* a conception of him, but of someone else?

What makes it a conception of him? There are basically two routes from the conception to Nagel. One is the route of *fit*. If Nagel's TN file includes enough detailed true stuff about Nagel, and nothing very important that's false, then he is probably if not the unique object denoted, at least the best overall fit for the conception. Is there a possible situation in which we have Nagel, with his conception and his thought, but the conception doesn't "pick out" Nagel --- he is not the best fit? This seems fairly straightforward. We just have to imagine a world in which Nagel is very different than he in fact is, and someone else is very much like Nagel in fact is. Of course Nagel's objective conception of TN may include some necessary truths; perhaps the objective conception includes who TN's parents are, are what his DNA is. But it seems like enough repeated applications of finding alternative denotations should eventually lead us to the desired situation, although its hard to be sure. However, I don't think *fit* is what makes the conception a conception of Nagel in the first place. It might be a conception of Nagel even if it not only might not fit him, but did not fit him.

The other route from the conception back to Nagel is from the conception to its *origin*. Nagel's impersonal conception, his TN notion, in my language,

belongs to a network of such notions. The network began when Nagel was born, or perhaps a bit before. His parents saw him at birth, they formed notions of him. They showed him to relatives and friends and neighbors, who formed notions of him. They sent telegrams and wrote letters to distant friends and relatives, perhaps enclosing photographs. They named him, creating a convention whereby people could use the name "Tom Nagel" to refer to him. A network of notions, in different people's minds, but linked to others by the purposeful exchange of information aimed at effecting one another's notion of the same person, was set in motion. It continues. I'm part of it, and so are you, since we are talking about Nagel in this essay. And so is he. Normally, he would be the one and only person whose self-buffer was also a part of this network, because normally a person can use their self-buffer as their permanent notion of themselves. But conception has gone temporarily impersonal.

Now suppose that Nagel thought that he was Gil Harman, to pick what one assumes is a somewhat unlikely example. Don't worry about how this happened. Nagel has a notion of Gil Harman, part of the Gil Harman network, started by Gil Harman's parents. For some reason Nagel has mistakenly come to think that he is Gil Harman. So he has linked his self-buffer with his Gil Harman notion. Asked his name, he says, "I'm Gil Harman." Asked his opinion about cognitive science, he says, "I think it is just terrific." And so forth. Now the point is, he would be wrong. When he said or thought, "I am Gil Harman" the "I" would refer to Tom Nagel and the "Gil Harman" would refer to Gil Harman. They wouldn't refer to the same thing, even though they were guided by the same internal notion.

Now what of his thought, the one he expresses with "I am Gil Harman." We might have here what I have elsewhere called a mess (Perry, 2001b). That is, one notion which is of two different people, with the information thoroughly mixed together. It is the same sort of thing that happens when you think you recognize someone as someone you have met before, but you are wrong, and it

takes you a long time to figure it out. However, let's assume that although Nagel has gone a bit crazy, he is a little worried, and keeps his identification with Harman tentative ---- the way I might have kept my identification of myself as John Perry tentative, until I had called my wife and checked that she recognized me. Think of a link between the self-buffer and the Harman concept, rather than an outright merger.

Even though it is tentative, his identification is false. What would it have taken for it to be true? At the subject matter level, there is no way it could have been true. Nagel is Nagel, Harman is Harman, and there is no possible world in which two things are one. But what possibility did Nagel confusedly have in mind? He thought that his impersonal conception of Gil Harman was of him; that is, that he was the origin of the network of which his impersonal conception of Gil Harman was a node. And that, I think, is a possibility. We can imagine worlds in which the very network that started with Gil Harman's birth started with Nagel's. We can imagine worlds in which the Sherlock Holmes network does not start with Conan Doyle at his study but with the birth of a child destined to be a detective on Baker Street

Given this exercise, it seems clear what we ought to say about Nagel's philosophical thought, "I might not be TN." The kernel of this, "I am not TN" is not a subject matter possibility. But it is a conceptual possibility, a possibility for how our thoughts fit into the world. What is at issue here is one of the alternative contents, what I elsewhere have called "network content."

At this point one may protest that this is certainly *not* what Nagel is supposing, when he supposes that he is not TN; he is not supposing something about the network that ties his and other notions of him to past talk about his and eventually his parents' decision to call him "Tom." But I agree that he is not supposing that. What Nagel is supposing is that he is not TN; he is supposing something impossible. What I am finding is a contingency that is not *what* he is supposing, but is a contingently false thing, such that if it were true his

supposition would be true. What he actually supposes would not be true, but, if this contingent falsity were true, his supposing would be a supposing of something else, which would be true. This contingent falsity is the condition that his impersonal TN conception does not have Nagel as its origin. That falsity is one of the contents of Nagel's thought, "I am not TN," and its contingency accounts for the phenomenology of the philosophical moment, the moment when the connection Nagel has with TN seems a matter of contingency, accident, and, I would have thought, in terms of choosing a philosopher one might want to be, good luck.

ⁱ I do not think Nagel intends to use "sense" as a theoretical term within a Fregean account of meaning, but simply as the way one is thinking of oneself when one uses "I" in this philosophical setting. At any rate, that's how I shall use it.

ⁱⁱ Nagel sets up a problem in section I and II of (Nagel, 1983). He outlines his solution in sections VI and VII, and there he introduces the "objective self." These sections are the core of his paper, insofar as we want to understand his positive view, and they are all I try to deal with. Sections III and IV are a critique of attempts to solve his problem without postulating an objective self. Sections VIII and IX qualify his views in subtle ways and if I dwell on them I find I have only a vague idea what is going on, so I am ignoring them. The version of this essay that appears as a chapter of Nagel's book *The View From Nowhere* (1986) has a number of changes.

ⁱⁱⁱ Sometimes it seems that there is meant to be only one objective self, rather than one to a person. This clearly won't help the present difficulty.

^{iv} With some exceptions, of course; I particularly recommend Perry, 2001b.