partial. Suppose, that is, that Rip Van Winkle and his neighbor Harold who has just awakened him both say "Yesterday was a fine day." I want to say that there is a belief state they are both in. Since they are both in it, it is a universal. And since they are doxastically dissimilar in many ways, it is partial.

One way in which Rip and Harold differ is beliefs about their context—that is, beliefs about the objects that are involved in the interpretations of indexicals and demonstratives they might use. Rip, we may suppose, believes that it is June 20, 1770, while Harold knows that it is June 20, 1790. These differing beliefs about the context will lead them to different behavior in certain circumstances. If we could persuade them to take a trivial true/false test, Rip would check "Today is June 20, 1790" false, while Harold would check it true. Their total belief states are not the same, but they still have something in common.

To determine what someone said, we consider the meaning of the sentence they used and the facts of the context in which they said it. To explain *why* they said what they did, we would need to consider the meaning of their sentence and their beliefs about their context. This is what will determine what they think they are saying. Consider a version of an example from David Kaplan. He is giving a lecture in a familiar seminar room; he gestures behind him where he believes there to be a picture of Rudolph Carnap, and says "That man is the greatest philosopher of the twentieth century." We will assume, for the sake of the example, that he believes that Carnap is the greatest philosopher of the twentieth century, that he wants to convey to the audience that he believes this and thinks saying it is the way to do that, that he thinks that "that man" refers, in a context, to that man who is most salient in the context, and that because of his gesture and the picture, he believes Carnap is the man most salient in the context. All of that explains why he speaks as he does. But if the picture has been changed without Kaplan's noticing it to one of Dan Quayle, he will not have *said* that Carnap was the greatest philosopher of the twentieth century, but merely tried to; he will have said something he did not intend to say.

This suggests that things do work in just the way I say they do not work in the paper. In the example, Kaplan has beliefs about his context and beliefs about English and beliefs about the history of philosophy. These beliefs explain his acceptance of "That man is the greatest philosopher of the twentieth century." But what I deny in the paper is that belief, conceived as a relation to context-independent sentences or Fregean thoughts, explained acceptance. When we press on this example, we will see that it does not count against that claim. Kaplan's beliefs about his context, for example, have to be what I call self-locating beliefs. He is not motivated to speak as he does merely because he believes that David Kaplan is in a position to express the proposition that Carnap is the greatest, but because he believes this in a certain way. He believes it by being in such states as the one I characterize as accepting "The picture behind me is of Rudolph Carnap."

Perry, Paper in The Problem of the Essential Indexical & Other Essays

4

A Problem About Continued Belief

I believe we have and probably need the notion of "continuing to believe the same thing." Suppose, for example, that Julius points to a man at a party, saying "That fellow near the bar is a dean." I believe him, saying the same sentence to myself. Later on, Julius is near the bar, the man is next to me, and I say "So you are a dean." A week later I see him in the quad and think to myself, "He is a dean." I think I came to believe something when Julius spoke, and continued to believe it until the end of the story.

Another example. My father told me "Santa Claus has a white beard." I continued to believe this for a couple of years, until I gave up believing in Santa Claus altogether.

In this paper, I first offer an analysis of continued belief in terms of an account of belief that distinguishes sharply between acceptance of a text and belief in a proposition. This account readily suggests two analyses of continued belief, neither of which seems quite right. Nevertheless, it seems to me that the account is supported on balance by its ability to distinguish between various possible concepts of continued belief, and so to finally suggest a less obvious but more adequate analysis. I then consider a problem, that of internal identity, that seems to arise with the suggested account. This problem leads to a theoretical notion, that of a file, which results from the adaptation of some ideas of Keith Donnellan's.

Texts

Suppose an articulate, sincere adult, eager to please, is asked to describe his beliefs. The result, we may suppose, is a set of true sentences of the form "I believe that S." The set of sentences S contained in these avowals are the *belief texts* from that subject at that time.[20] They seem to tell us something important about the believer. For example, if one of them were "I will be killed if I do not

[20] I borrow the term "text" from Michael Pendlebury (1982). He got the term from Gustav Bergman.

leave the room immediately," we might reasonably expect the subject to try to leave the room. Now we can assign even obstinate inarticulate adults sets of belief texts, in virtue of the sentences of the form "I believe that S" that would be true, if they were clever enough to think of them and cooperative enough to produce them. Perhaps if we include in the set of possible texts sentences or groups of symbols not in ordinary languages, consisting, say, of pictures, we might extend the notion to include children and even less articulate believers. At any rate, whatever problems are involved in the notion of a text, it seems to me to be one that is implicitly used in virtually all attempts to philosophize about belief, and so it is of some value to get it out in the open and think about it, as I try to do in this paper. (See also Essay 3.)

What is Believed

I think that belief, like perception, is a relation we have to things not in our heads, partly in virtue of what goes on in our heads. If, as I shall put it, I *accept* the text "My wife has brown hair," then I am related to two things that are not inside my head, my wife and the property of having brown hair. Corresponding to these two things are a number of objects. For certain purposes, I think we need to distinguish:

1. the situation that makes it the case that she has brown hair;
2. the fact that she has brown hair;
3. the state of affairs that she does have brown hair, and the state of affairs that she does not have brown hair;
4. the propositions that she does have brown hair, and the proposition that she doesn't.

I think situations are the most concrete objects, and objects and properties are uniformities across them. States of affairs are abstract objects that situations make factual or not factual. States of affairs give us the basic propositions, but propositions also include more complex conditions such as, for example, there being more than one philosopher who has owned a 1939 Pontiac.

The issues in this paper do not depend on accepting all of these objects and distinctions, however. I am concerned with belief in propositions of a simple kind I call *R-propositions*. These can be identified by a sequence of an n-ary relation and n objects of the appropriate sorts. A proposition $\langle R; a_1, \ldots, a_n \rangle$ is true if $a_1, \ldots, a_n$ stand in the relation R, false otherwise. R-propositions are what Kaplan calls *singular propositions* (1978, 16). As he points out, sentences using indexicals and demonstratives seem to express R-propositions. If I say, "I wrote this paper," and you say to me, pointing at the same thing, "You wrote that paper," we have said the same thing. We have both expressed the same

R-proposition, even though we identified the individuals in it differently.[21]

I use the term *R-proposition* as short for *Russellian proposition*. The identity of these propositions depends on the identity of the individual in them, not on the way it is identified. (Has brown hair; my wife) is the same as (Has brown hair; Frenchie Perry). This means that the proposition does not contain the *mode of presentation* or *cognitive fix* that the believer has on the individual. Russell thought that these sorts of propositions were expressed by sentences with real names in them. Still, it might be misleading to call these Russellian propositions. He did not think there were very many real names, and in the history of the philosophy of language he is most famous for his theory of descriptions, which is a way of avoiding R-propositions. A sentence with a definite description in it will not express an R-proposition, according to Russell. It will express a proposition that is true if there is someone who satisfies the description and has the predicated property. Propositions of this latter kind are close in spirit to what Frege called *thoughts* (although the theoretical apparatus of Frege and Russell differs a lot). Frege thought that R-propositions were a rather weird idea. (See Essay 13 for more discussion of these issues.)

Role and Proposition

Once my older brother told me "I will give you five dollars tomorrow." I believed him, and asked him for money the next day. He said, "I said I will give it to you tomorrow, and I will." This kept up for a couple of days until I got the message. At the time, I thought the message was that tomorrow never comes. Now I realize that it does, though when it does, we do not call it "tomorrow." The real point my brother was trying to get across was this: Do not confuse sameness of text accepted (or asserted) with sameness of R-proposition believed (or asserted).

Each day my brother induced me to accept the same text: "He will give me five dollars tomorrow." By accepting this, each day I believed something dif-

[21] In Essays 1 and 2, I argued that we need to adopt some such notion to handle thinking and speaking that makes use of one's identity and position in the world. These were called "thoughts" in the first paper and "objects of belief" in the second. But the nonindividual constituents of thoughts or objects of belief were not properties. In the first paper, they were to be Fregean senses of predicates, and in the second paper, the matter was left obscure. It now seems clear to me that we want properties and relations, not senses or other kind of meanings as constituents of what is believed. We may say that situations are composed of references and not senses, thus describing in Fregean terms a notion Frege did not adopt. Here it must be remembered that Frege took properties—not extensions—to be the references of predicates, and that properties are distinguished by Frege from the senses of predicates. (See Frege 1891/1960 and 1892/1960a.) (There is a tendency on the part of Fregeans to let senses of predicates do the work traditionally done by properties, and move extensions into the position of the reference of predicates. This tendency has a basis in Frege's own practice for his concepts may be thought to differ from extensions—more precisely, courses of values—only in being unsaturated. But in fact, he did distinguish concepts from senses and from extensions, and did call them "properties" on occasion.)

ferent. On Sunday, I believed (x will give y $5 on z, my brother, me, Monday); on Monday, (x will give y $5 on z; my brother, me, Tuesday); and so forth.

So we can believe different R-propositions in different contexts, in virtue of accepting the same text. We can also believe the same R-proposition, in the same or different contexts, by accepting different texts. I may believe (x has brown hair; my wife) on a given occasion both in virtue of accepting "my wife has brown hair" and "the person on my left has brown hair," and I might do so whether or not I accept "the person on my left is my wife." And you may also believe this very same object, in virtue of accepting yet another text, say "His wife has brown hair."

Behind this picture of belief are some semantical notions, which I want to make a bit more explicit. Expressions are thought to have *roles*, which are taken to be an important aspect of their linguistic meaning. A role is a rule that takes us from a context to a reference. A context is identified by a person and a time—the speaker or thinker and the time of the thinking or speaking. Thus the meaning of "I" is given by the rule that it stands, each time it is used in thinking or speaking, for the thinker or speaker. Given the context—say, President Carter on some given day—the rule gives us the referent of a use of "I": President Carter. "Is red" has a role that takes us from any context to a certain color property, while "is the color of the typewriter I am now using" takes us to different color properties from different contexts. A simple atomic sentence, like "I am sitting," takes us to an R-proposition.

This semantical picture is an embodiment of some ideas of David Kaplan, in his extremely important studies of the logic of context-sensitive expressions (1978 and 1979). But it differs from his system of *characters* and *contents* in crucial ways, and since "character" and "content" are precisely defined notions, it seemed best to adopt some new terms. Kaplan's characters are functions from contexts to contents, and so start off very much like roles. The important difference comes when we get to the contents. Kaplan's contents are *intensions*. They are explications of, or at least descendants of, Frege's absolute senses.[22] Thus for Kaplan "I" takes us from a context with me as speaker to a constant individual concept of me, an intension that takes me as extension in each possible world (1979, 403). A role, however, takes us directly to me, with no intervention of an intension.

Now these differences may be simply terminological[23] and in any case, what will be said about continued belief does not turn on the difference between role and character. But I suspect the difference is more than termino-

[22] Frege's senses were absolute in the sense that their reference was not relative to a speaker or context.

[23] I say this because Kaplan 1975 suggests adopting a Russellian approach to intentional logic. I am not sure where this would leave the system in Kaplan 1979, but it seems it might lessen the differences just mentioned.

logical, and so would like to briefly explain my revisionism. It seems to me that Kaplan has not made radical enough changes in the Fregean system; like many who hit upon revolutionary ideas, he hesitates at completely uprooting the structures within which he grew up. Kaplan sees character as an aspect of sense that Frege missed, something that needs to be added on top of sense. So in his system we have, for each possible use of an expression, its character, its content in the given context, and its reference, determined by the content. Contents, like Frege's senses, will not determine references relative to a context, but absolutely.

But I think Frege's notion of absolute sense was a *mistake*. We need something like character *instead of* absolute senses or intensions. By "absolute" I mean that Frege's senses have a reference independent of context. My picture is this. Frege introduces senses, in the first instance, to serve as ways of apprehending references. We need the notion of a way of apprehending a reference, but Frege's notion of sense will not do, because absolute. He cannot bring out the difference between the way you think of me when you think "you" and the way I think of me when I think "I." That is because these ways of thinking do not correspond neatly to what is thought about. Ways of thinking cut across what is thought about. The way I think of myself when I think "I" is just the way you think of yourself when you think "I." Only the way of thinking plus a context determines a reference. And that means we need something like character or role, something that takes us from the subject's identity and position in the world to reference, to serve as ways of thinking.

However, because he took senses to be absolute, Frege was free to use them in ways in which characters or roles could not be used. The additional uses for senses become an integral part of Frege's system, so we cannot make his senses nonabsolute without altering it in important ways. The relation *sense of* provides a homomorphism from the realm of sense into the world of things and properties, allowing the former realm to do much of the semantical work of the latter. Belief is a case in point. That one may believe that a is a featherless biped without believing that a is human, even if the featherless bipeds are just the humans, was explained by Russell in terms of two different properties having the same extension. But in Frege's system, it is explained by the different senses of "human" and "featherless biped." We no longer need to address the question whether different properties can have the same extension. Fregeans cease to distinguish it from a question about expressions and senses. Most importantly, the system of absolute senses allows Frege to take sentences to stand for truth-values, and ignore R-propositions altogether. There are properties in Frege's system—they just are not given much to do. But R-propositions are not even part of the system. It does not matter that in the reference of the sentence "all that is specific is lost," as long as the specificity is retained in the sense of the sentence.

I think that Kaplan has shown that absolute senses cannot by themselves handle the ubiquitous phenomena of context-dependence. I think attention to the epistemology of context-dependence suggests that the phenomena are not only ubiquitous but ineliminable. So senses cannot do what they were introduced to do—serve as ways of believing. And I think that, by recognizing properties and R-propositions, the remaining work for absolute senses is removed. We can let the objective world handle the work of absolute senses. (See Barwise and Perry 1981, 1983.)

Belief States

While I have spoken, and will continue to speak, of believing an object in virtue of accepting a text, this is misleading. I think texts accepted are of interest as a way of characterizing *belief states*. One may think of belief states as brain states, perturbations in mind-stuff, or whatever—right now I do not care.

My point is that *texts accepted* is a useful notion for classifying believers, when we want to look for plausible general principles with which to explain and predict human behavior, and that this is clearly because roles correspond to internal states. My own view is that belief states are multiple manifested dispositions, and that with the notion of a text we label the state in virtue of one of its manifestations. But in any case, I want to emphasize that I do not think of texts or their roles as "direct" objects of belief. They are not, ordinarily, objects of belief at all, but objects produced by one who is in a belief state, under certain conditions.

So, strictly speaking, it is not partly in virtue of text accepted, but partly in virtue of one's belief states, that one believes the objects one does. The other factors are positional and causal factors.

In the case of texts containing indexicals, this basically means the context of the belief: the identity of the believer and the time at which the believing occurs.

In the case of texts involving proper names—and, as we shall see, other cases—it includes the causal history of one's use of the name. Suppose there are two English speakers whose texts include "David Lewis has a beard." There is a belief state they both are in. Given that each hated anyone with a beard, both would be unhappy if told "David Lewis is here to see you." Nevertheless, on the assumption that David Lewis the philosopher and David Lewis the railway magnate are different men, their beliefs are about different people.

So, in virtue of being in a belief state in a certain environment, we believe a certain object. Because the same object can be believed in different ways, from different environments or "points of view," classifying people by the objects believed is not always particularly useful.

Both you and I may believe (x will get hit if he does not duck; John Perry).

I duck, you shout. Why the difference in behavior, if the same object is believed? Because what is important is *how* it is believed. I accept "I will be hit if I do not duck," you accept "He will be hit if he does not duck." But I act just as you would if you accepted what I did.

There is a traditional view according to which what one does is explained in part by what one believes. Now the view I am putting forward is different. What one accepts explains in part what one does; what one accepts plus the context in which one accepts it, determines what one believes. This can seem very puzzling. If what one does does not depend on what one believes about the world, why expect it to be appropriate to one's situation in the world? There seems to be a sort of preestablished harmony involved. Another way of approaching the same question is to note that the accepted sentences really have two roles. What I have been calling their role is their *semantical* role. But they also seem to have a *cognitive* role. How do the two connect?

While I cannot here develop an answer to this question, I want briefly to suggest one. Consider the sentence "There is a hungry lion coming towards me." Now consider the contexts relative to which this sentence is true. They all consist of persons and times such that the person is being approached by a hungry lion at the time. It is a good idea for all of these people to run like crazy. In a sense they do not need to know what they believe. Even if they have forgotten who they are and lost track of time, they know enough to run. (In another sense, they do know what they believe even then.) Most of these people will not believe the same thing. But each of them will believe in *something* that provides *them* with good reason to run.

Kaplan develops the idea of a "theorem in the logic of demonstratives." This is a character that, though it may lead to different contents from different contexts, and may not lead to anything necessarily true, always leads to a content that is in fact true. His favorite example is "I am here, now."

I am relying on a notion with a similar structure. "If I am approached by a hungry lion, I should run like crazy" is *reliable advice*. Each person who accepts it will be well served. But in a sense, each person who accepts it is not getting the same advice. I am adopting the strategy of running when I am approached by a hungry lion, while you are adopting the strategy of running when you are approached by a hungry lion, And the advice I am adopting *would not* be good for everyone to follow. Why should you run when I am approached by a hungry lion?

This all suggests a way in which the cognitive role and the semantical role of a sentence are related. In a sense it is a preestablished harmony. But it is not a mysterious preestablished harmony.

Preservation of R-Proposition Believed and Continuity of Belief

I will now return to the topic of continuity of belief. The picture of belief sketched suggests two obvious conditions for it. The most natural, perhaps, is that it consists in preservation of the R-proposition believed. A second and different hypothesis is that continued belief is simply retention of text accepted (that is, retention of the belief-state that is a disposition to accept a certain text). It seems our ordinary notion of continued belief might be either one of these notions, or some combination thereof.

Variations on the Morning Star/Evening Star case seem to show that preservation of R-proposition believed is not a sufficient condition of continued belief. Suppose I ask two friendly Chicagoans which building is Union Station. One points to a building the three of us are standing beside, the other to a building some distance away, across a street. Both are correct, as Union Station is mainly under the street, rising up on either side of it. But I do not know this. My mind changes as I turn to each of the honest-looking natives. First I accept "This is Union Station," and then "That is Union Station." The R-proposition believed remains constant: (x is Union Station; Union Station). But this does not seem to be a case of what we would ordinarily call continuing to believe the same thing.

We might think that the case of continued belief in Santa Claus' white beard shows that preservation of R-proposition believed is not a necessary condition either, for here we have continued belief with *no* R-proposition believed. But this is pretty indecisive. We might say instead that the R-proposition is *nothing*, and is preserved, or perhaps that the R-proposition is (x has a white beard; the null set). If we had good general reasons for thinking that preservation of R-proposition was a necessary condition of continuity of belief, we would try to handle this sort of case one way or another.

A more convincing case comes from Pendlebury (1980, 1982).[24] Someone being driven around the environs of Bloomington, Indiana, retains the text "This is Monroe county." Unknown to her, she leaves and reenters Monroe County several times in the course of the drive. So the R-proposition believed changes. But isn't this a case of continued belief?

I find that different people say different things about this example. Philosophers (with exceptions, including Pendlebury) tend to think that as she moves in and out of the county, she acquires new beliefs. Nonphilosophers (e.g., my wife), whom we may think of as having purer linguistic intuitions than philosophers, tend to agree with Pendlebury, that the person in the example continues to believe the same thing, although what she believes changes truth-value. For the moment, I will simply leave this disagreement, and the question of whether preservation of object believed is a necessary condition of continued belief, unresolved.

[24] I read Pendlebury's paper and discussed these issues with him while writing this paper and this was very helpful and illuminating. Pendlebury uses this example as an argument against the notion of an object of belief, and develops an account of belief that takes belief to be mainly (in the terminology I am using) a matter of what goes on inside, with no important role for external objects of belief.

Retention of Text and Continuity of Belief

The results of the last section suggest that continued belief does not have just to do with the outside factors in belief, so we should look inside. But the simplest hypothesis, that continued belief is just continued acceptance of the same text, is quite wrong. Retention of text is neither a sufficient nor a necessary condition of continued belief.

My brother showed that it is not sufficient. I continued to accept "He will give me five dollars tomorrow," but did not thereby continue to believe what I first believed. Another example: Suppose I am tied to a chair, waiting for the time bomb on the table to explode. At each time I accept "It is going to explode an instant from now." This is not continued belief, but continuously changing belief.

Nor is it necessary. The case described at the beginning of the paper, where I believed a certain man to be a dean, is one of continuous belief. Yet no text is retained. First we have "That man is a dean" and then "You are a dean" and then "He is a dean."

Preservation of R-Proposition by Retention of Text

Suppose Julius had said, "That man next to the bar is Halsey, and Halsey is a dean." I retain the text "Halsey is a dean" throughout the party, and up to the day I meet him on the quad. Surely this is a paradigm case of continuing to believe something. And it suggests at least a sufficient condition: we continue to believe when we continue to believe a given R-proposition by continuing to accept the same text.

If this suggestion were correct, it would be tempting to try to use it to reply to the last claim of the last section, that retention of a text is not a necessary condition of continued belief. After all, that I thought of Halsey in different ways during the period of time in question—as "that man" and "you" and "he"—does not show that I did not also think of him in some other way the whole time. *If* Julius had told me his name, I would have thought of him as "Halsey" throughout, even though I also thought of him as "that man" and "you" and "he." Yielding completely to this idea, we might argue as follows. Clearly what we call continued belief is just preserving the R-proposition believed by retaining a text. Since in the original Halsey case (where I did not have a name for him) I continued to believe, we know that there must be a text

I continued to accept. I just did not look hard enough for it.

Tempting as this is, I do not think it can be correct, if to have such a text I must have some way of thinking of or referring to Halsey of the ordinary sort, which I can use throughout the period I believe him to be a dean. I would have to possess some singular term *A* that designates Halsey all along. But what could *A* be?

A might be a context-insensitive definite description that denotes Halsey. Indeed, the meaning of such a thing is just what would have been taken, combined with the meaning of "is a dean," to compose a traditional proposition. Such propositions, on the traditional theory, serve the purposes of both my texts and belief objects: they are the bearers of truth and the classifiers of belief states.

But I see no reason whatsoever to suppose that I could produce, or have in my possession in any manner, however implicit, such a definite description for Halsey (or for that matter for any other object).[25]

For the purposes of the temptation at hand, however, such a context-insensitive definite description is not needed. A singular term, like a context-sensitive definite description, or an indexical, or a proper name, would do as well. The designation of *A* need not be Halsey from anywhere at any time, just from those places I occupy between the party and the day on the quad. "The man on my left" would do, if Halsey happened to be the unique individual on my left throughout the episode.

Even so, I cannot see that I need to have such a singular term in my repertoire to continue to believe what Julius told me. Someone might argue that my problem is an insistence that the singular term be one that can appear in a public text—that perhaps an "inner name" of some sort is being employed instead. Little as I understand it, I have no objection to introspection as a pastime or even as a method for doing philosophy. I just can find no inner name that I would have to have or be likely to have : the sort of R-proposition described.

In any case, I do not think preservation of R-proposition believed by retention of text accepted is even a sufficient condition of continued belief. It is almost sufficient. When we see what else is required, we will have a clue for handling cases in which a single text is not retained.

The example I have in mind to show nonsufficiency is yet another variation on the morning Star/Evening Star case. Smith, whose watch is an hour fast, accepts "Today is my husband's birthday." Just very shortly before eleven, she looks at the calendar and realizes that she had it wrong. It is March 1 and not March 2. Just as this sinks in, she glances at her watch—precisely at eleven, so it shows precisely midnight—and thinks to herself, "so *today* is my husband's birthday." And by doing that she preserves the R-proposition believed, for in

[25] I now think this is much too strong. See the last section of Essay 13.

this case she does believe {*x* is my husband's birthday; March 1} even though she also believes {*x* is my husband's birthday; March 2} and even though she does not accept "March 1 is my husband's birthday." So she preserves the R-proposition of belief by retaining a text. But she does not continue to believe the same thing. She changed her mind.

Retaining a Text in Order to Preserve R-Proposition Believed

The problem here is that Smith was not trying to preserve object believed: indeed, she was trying not to. This suggests the following view.

We have a (fairly) clear notion of preserving R-proposition believed. This is an "external" notion, having to do with the R-proposition believed, however it may be believed. But we have a device for generating an internal notion from the external one. We suppose that there is such a thing as *trying* to preserve R-proposition believed—or, to diminish the suggestion that conscious effort need be involved—retaining one's texts *in order to* preserve R-proposition believed. This is an internal notion, for one can retain one's texts in order to preserve R-proposition believed, but not succeed. Let us call this internal notion *internal continuity of belief*.

Now perhaps we really have two ordinary notions of "continuing to believe." One is just internal continuity of belief. This is the one Pendlebury and my wife employed, in their reaction to the episode in (and out of) Monroe County. The other is *successful internal continuity of belief*. This suggestion won't quite do, however, for it implies that retention of text is a necessary condition of continuing to believe the same thing, but we have already seen that it is not.

Change in Text in Order to Preserve the R-Proposition Believed

Consider the original case, which was used to show that retention of text is not a necessary condition of continuity of belief. I first accept "That man near the bar is a dean" and then, after Halsey has left the bar and walked over to me, "You are a dean." It seems to me this is not only a clear-cut case of preserving the object believed, but also a case of internal continuity of belief. This is what it is like to try to continue to preserve the R-proposition believed, when one is aware that a constituent of the R-proposition believed is moving around.

So it seems that not even our internal notion can be identified simply with retaining a text. Rather, the general case of continued internal belief is changing the text in order to preserve the R-proposition believed, where we understand that change includes retention as a limiting case.

So now we have three notions:

(i) preservation of R-proposition believed, or external continued belief;

(ii) changing (or, in the limiting case, retaining) the text *in order to* preserve R-proposition believed (internal continued belief);

(iii) successful internal continued belief (ordinary continued belief).

The Problem of Internal Identity

Now I finally reach the problem of the title. I have said that we have the notion of *internal continued belief*. Thus far, I have assumed that the internal states are displayed by the structure of texts accepted. It seems then that continued internal belief should be displayed by relationships among the texts accepted. But saying how it is displayed is a problem, one that will lead us to a new theoretical notion, that of a file.

In the example concerning Halsey, I used "you" at the later time in order to refer to the same person I had used "that man near the bar" to refer to earlier. I shall call this a case of *internal identity*. Internal continued belief in the Halsey case requires internal identity and what we might call internal sameness of attribution. I ignore the conditions for the latter, and spend the remainder of the paper on internal identity.

We might suppose that internal identity would be displayed by identity texts. But neither at the earlier or the later time do we find me accepting "You are that man near the bar." At both times, this would have involved me in the belief that Julius and Halsey are one, which I never had.

At the later time we would have found "You are the man I referred to as 'that man near the bar' a moment ago." The presence of this text does seem to show, in a sense, that at the later time I am trying to refer with "you" to the person I earlier referred to with "that man near the bar." But it is not the right sense for internal identity.

Suppose Lew was near the bookcase while I was talking to Julius, and Julius pointed to him saying, "That man near the bookcase is a historian." At the later time, while talking to Halsey and pointing to Lew, I accept these texts: "You are a dean, and were standing by the bookcase a moment ago" and "He is a historian and was standing near the bar a moment ago." Both texts are false, for I have become confused. Perhaps there are different ways I could have gotten confused, but one way seems to be this. I come to believe and continue to believe both {x is a dean; Halsey} and {x is a historian; Lew}. But I misremember where each man had been standing earlier.

In this case, I would accept "You are the man I referred to as 'the man near the bookcase' a moment ago." This would be false, of course, but that is not my point. I would also not display internal identity. Just because at the end of the later time I accept "*A* is the person I referred to as *B* a moment ago," it does not follow that my earlier use of *B* and later use of *A* constitute internal identity. In this case, my earlier use of "that man near the bar" and later use

of "you" are a case of internal identity, even though I accept "You are not the man I referred to as 'the man near the bar' a moment ago."

I do not think that the solution to the problem of internal identity involves finding identities among the texts. And even if, contrary to what was claimed earlier, we could find, whenever we have a case of continued belief, a retained text, our problem would not be solved. Suppose we had found the mythical *A*, the singular term with which I referred to Halsey throughout the episode. Then we might suppose we could reconstruct the process of internal continued belief as a sort of inference:

(Earlier time) That man near the bar is a dean.
That man near the bar is *A*.
So, *A* is a dean.

(Later time) *A* is a dean.
You are *A*.
So, you are a dean.

But now we are assuming that the use of *A* at the earlier and later times is an instance of internal identity. But we cannot assume that, as was shown by the case of Smith and the birthday.

Internal Identity Without Singular Terms

Moreover, another kind of case will show that any reliance on identity of singular terms from the earlier to the later time, or identities formed with singular terms formed at the earlier or later times, must be misguided. For we can have internal continued belief and internal identity when I have no singular term at all to refer to the thing I continue to think about.

Consider the period of time after the party, and before I saw Halsey again on the quad. I think it is perfectly possible that I continue to believe (x is a dean; Halsey) during the interim, even though I have no way of referring to him at all. He is indexically inaccessible—"you" and "he" and "that man" will not reach him. I know nothing about him that does not apply equally well to many others. I might, of course, remember enough about the circumstances to think of him with a description like "The man who was near the bar while I was talking to Julius," but I need not remember this, so far as I can see.

In order to make the phenomenon of recognition intelligible, we shall have to suppose that I believe more about Halsey than merely that he is a dean. Let us suppose that, as Julius directs my attention towards Halsey, I come to accept "That man near the bar is a dean, wears a bolo tie and a rumpled suit, smokes a pipe, and has a mathematical demeanor." Seeing a man on the quad with a pipe, and this sort of suit, tie, and demeanor, I say "Ah, yes. You are a dean." If this is the man Julius pointed at, then we have a case of recognition.

The information about Halsey seems to have been retained in my mind without benefit of singular terms. So it cannot be the singular terms that are the key to internal identity. It seems, instead, to be just what was retained in this case, the cluster of predicates, that provides the internal identity.

Predicates do not need a singular term to be knitted together. A quantifier phrase serves this purpose: "Some man wears bolo ties, rumpled suits, smokes a pipe, has a mathematical demeanor, and is a dean." Such a sentence among my texts we would naturally think of as that in virtue of which a general belief would be held, rather than a belief about a particular person. Yet this seems a mistake (or perhaps better the phrase I use next can be considered a singular term). Consider the phrase "A certain man"—described in the dictionary as meaning "not further described or more specifically named but assumed to be known." If I say "A certain man wears a bolo tie and rumpled suits," I would be ordinarily thought to have someone in mind and not merely to be making a likely guess that at least one person in the world dresses in this way. "A certain man" functions here, I shall say, as an ersatz name.[26] In virtue of accepting the sentence, I believe of some specific individual that he dresses in the manner described. But how does this man get involved in the belief?

Files

I believe the solution to this problem lies in the adaptation of some ideas from Keith Donnellan's account of proper names (1970 and 1974). This adaptation leads to a new theoretical notion, that of a file.

On Donnellan's account, a proper name refers to an entity on a given occasion of use if that entity occupies a unique position in the correct historical account of that use. Suppose, for example, a student in a seminar tells me her name is "Phyllis." Later on, in talking to a colleague, I say "Phyllis has studied Frege's writings with great care." A complete account of the factors that led up to this utterance would mention many people, including me and my colleague. But it seems that the student would figure into the account in a certain way that others would not. That is what makes her the referent of my use of "Phyllis."

Now it seems to me to be helpful to speak of causal chains here (though I doubt that it would seem so to Donnellan). The property is having a certain role in a causal chain that led up to my use of "Phyllis." One may object to this account because it is vague, for Donnellan says very little about what this role is. This seems to me not a good objection, but simply a suggestion for further research. Suppose there were a debate about whether "is red" stands for a color property or a shape property. Close study of a number of examples chosen to minimize color/shape correlation would convince us that it stood

[26] The view that an indefinite description like "a man" can function semantically as a referring expression is defended in Wilson 1978 and in Chastain 1975.

for a color property. This would be quite clear, even while we remained quite unable to say much about which color property it was, other than by pointing out instances of it. Examples Donnellan provides convince me that for any proper name A and any person Q, "Q refers to x with his use of A" stands for a historical or causal property, and not the property of satisfying a set of descriptions held by Q or even Q's linguistic community.

It may seem that Donnellan's account cannot be adapted for ersatz names, for the history of my uses of "a certain man" would not be of much interest in determining whom I had in mind on a given occasion. That phrase has nothing in particular to do with Halsey. But as Donnellan emphasizes, it is not the history of the use of the name, but the events that lead to a given use of the name, that are important (1974, 19). Donnellan's theory is compatible with a name being used for the first time in history referring to a long-dead individual. In the example, my use of "Phyllis" might be the first time the student had ever been referred to by a use of that name—if, for example, her name was "Philo" and I had simply heard it wrong. It might also be the last, if my colleague says, "You must be referring to Philo—there is no one named 'Phyllis' in your seminar." That I refer to Philo as "Phyllis" on this occasion does not make it her name, of course.

Thus we can apply Donnellan's account to ersatz names. Doing so leads us to see that it is not the name itself (ersatz or proper) that is of interest, but the grouping of predicates. It is the predicates grouped that lie at the end of a causal chain originating with the person I am thinking about.

Suppose then that on Wednesday I accept "Halsey wears a rumpled suit and a bolo tie and has a mathematical demeanor." By Friday, I have totally forgotten the name of the dean, but still accept "A certain man wears a rumpled suit and a bolo tie and has a mathematical demeanor." If the earlier acceptance and the later acceptance are parts of the same causal chain of the right sort, then we have internal identity. And if the chain leads in the right way back to a unique individual, that is the individual I think about on both occasions.

Of course, there are many ways we might have internal identity without there being a single person, in any clear sense, thought about. There might be no individual initiating the chain in the right way. Causal chains might merge, so that a week or so after meeting Halsey and Lew I accept "A certain man named Lew—or maybe Halsey—wears a bolo tie and is a historian." Although the facts are clear enough, there will no longer be a straightforward answer to the question "Who is this belief about?"

Let us say that the set of sentences a person accepts at a given time is their doxastic profile at that time. We are now supposing, then, that linking various entries in the doxastic profiles of a given person at various times are causal chains, and the entries so linked we shall say belong to a single *file*. Files are clearly theoretical notions, even relative to profiles, internal identity

is not *displayed* in the profiles. In this respect, this adaptation of Donnellan's ideas deviates significantly from the kinds of R-proposition he deals with. In his examples, it seems clear that we would know what to say about what a given use of a name refers to, given the sort of information we can readily imagine having, about who first said what, point at whom, and the like. But here we do not have this feeling. Given all the information about the goings-on in someone's brain for a period of time, I would have no idea what to say about whether two uses of a name were due to an entry in a single file. This seems to be a case in which we have mastered all sorts of indirect tests for a phenomenon without having the foggiest idea what we are testing for. But this is not an unusual situation. I know how to test to see if my television is tuned to the same channel as last night—even if the channel selector knob is slipping and the numbers on it are no sure guide, I can figure it out pretty well. But I have only the foggiest idea what being tuned to the same channel really is, and would be quite helpless if confronted by the inner works of a television and asked to explain the conditions under which we have the same channel tuned in.

Although I have no idea how a system of files might be instantiated in a brain, I can give an analogy that shows how such a system could be instantiated with a much more primitive information-storage system. Let us suppose that on the first day of class I carefully note down features of the various students around the table. I use full sentences, for example, "The student in the seat first to my left is a woman, has blonde hair, is short, was born in Ottumwa, is a psychology major, and has on a red sweater." I call this "opening a file." I do not use names (perhaps for religious reasons). Later in the day I look at the cards. I erase or alter some phrases. I erase "the student in the first seat to my left" and change "has on a red sweater" to "once wore a red sweater." At the next lecture, I use the remaining predicates for purposes of recognition. For example, to the student first on my right I say "Weren't you born in Ottumwa?" for she is a woman, blonde, short, with an *Advanced Psychology* text at her side.

There are no singular terms left on my card. The predicates remaining might allow me to construct one, perhaps "The only student ever to wear a red sweater and be a psychology major and be from Ottumwa." But I need not do this; it serves no purpose, and it would not change things if the description did not fit the student for whom I opened the file (maybe I heard wrong and it was Omaha, not Ottumwa). Although I could have done this in a more orderly way, assigning each student a number, or using their names, or assigning them a name of my own, I need not. The work that would be done with these devices is done by the predicates being written on a single card.

Now each such file card offers me, at each seminar, a profile, a set of predicates. Whom does this profile stand for? We clearly must make a distinction between the persons, if any, of whom the profile is true, and the source of the

profile. The source is the student my perception of whom led to the establishment of the file (it need not have been perception, though, as far as I can see). The predicates on the file may or may not be true of her. I update the cards after each class. Old entries are changed, new entries are made, the writing fades, I may have made mistakes recording the data, and so forth. I have come to believe something false about a certain student—a student I may no longer have any way of recognizing.

Instead of writing on a file card, I might have taken a photograph. The distinction between whom a photograph is *of* and whom it might resemble is familiar. (See Kaplan 1969.)

I think our notion of internal identity is based on thinking of the mind as working in this way. In the analogy, if the same file card has "is F" on it at one time and "is G" at another, then according to my files a certain person was F at the earlier time and G at the later time. This is the analogue of internal identity.

If we think of the mind as a storehouse or a filing cabinet, the relation would be between the stored items that correspond to predicates, say, being stored in the same room or in the same file drawer or folder. What is essential in these metaphors and analogies is a path from the production of texts at one time back to the original perception of (or other introduction to) the source at an earlier time. This path in the mind plays the role of an object in the world. So our notion of internal identity, and so ultimately of believing the same thing, depends on the identity of the internal causal path or chain.

Now it may be that in the mind the work of files are done by something like proper names. It is important to stress that identity of proper names does not constitute internal identity, and so we need the notion of a file to explain what would be special about these names.

One advantage of filing with proper names is the ease with which one can handle relational predicates. In the analogy I have presented, there is really no provision for handling such predicates. I'd like to end by pointing out how considerable relational information can nevertheless be handled in such a system, for this is a point that seems importantly related to the study of what it is to think from a position in the world.

First, the order of the files might be significant. Thus, in the seminar case, I might keep the cards stacked in the order in which the students are seated throughout each seminar. Second, context-dependency itself provides a way of embodying relational information in a system of one-place predicates. Consider the fact that a student wears a red shirt on a certain day. This is a relational fact, involving the student, the day, and the relation "x wears a red shirt on y." But on that day I can believe that fact simply by accepting "α is wearing a shirt." The fact that the student is to *my* left, *I* can believe simply by accepting "α is to the left." I do not need a file for the day, or for myself; I finesse the

need for such files by setting up the system on the day, and by being me.

We may think of the final product, at the end of the quarter, as a pack of bundles of ten cards each, each bundle giving information about what has been taken to be one student, the order of the cards in each bundle showing what I observed the first week, second week, and so forth. Through exploitation of context, and arrangement of the files, considerable relational information is embodied, without using names or other than one-place predicates. There are, of course, a lot more efficient ways to encode such information, but that we can, by using our position in space and time get by with monadic predicates seems to me to be a fact that may be of some importance, both in understanding thinking as a phenomenon that exploits our position in the world, and in understanding how such thinking can lead to increasingly objective conceptions of the world.

Afterword

The first three essays do not say much about what is involved in continuing to believe the same thing. When I was working on the first two essays, I assumed that belief consisted in (what is here called) accepting a text in a context. I was inclined to think that continuing to believe the same thing must involve adjusting the text accepted to suit change of context so as to preserve what is believed. This view is suggested by Frege's remark:

> If someone wants to say the same today as he expressed yesterday using the word "today," he must replace this word with "yesterday" (1918/1967, 24).

In "Demonstratives," David Kaplan brought up this issue under the heading of "cognitive dynamics":

> Suppose that yesterday you said, and believed it, "It is a nice day today." What does it mean to say, today, that you have retained *that* belief? It seems unsatisfactory to just believe the same content under any old character—where is the retention? You *can't* believe that content under the same character. Is there some obvious standard adjustment to make to the character, for example, replacing *today* with *yesterday*? If so, then a person like Rip Van Winkle, who loses track of time, can't retain any such beliefs. This seems strange (1989, 537–38).

He elaborated on this theme in some lectures he gave at Stanford in fall quarter 1978, discussing several examples of conversations involving continued reference to individuals who had become indexically inaccessible. These brought up the problems of "internal identity" very forcefully. This paper was written for the Oberlin Colloquium in winter 1979. I am grateful to many participants,

particularly John Perry, for comments.

In the original version, I eschewed the notion of a proposition, taking what is believed to be "situations." The notion of a situation I had in mind was essentially what Jon Barwise and I called "abstract situations" in *Situations and Attitudes*. In the present version, I have instead used the terminology and notation of the later essays in this volume, so that what is believed, in the cases under consideration, is an "R-proposition." I have come to think that both the motivation Barwise and I have for situations, and the difficulties our approach presented with respect to propositions, are not directly relevant to the main point of this paper.