

Rip Van Winkle and Other Characters*

John Perry[†]

March 4, 2009

‘If someone wants to say the same today as he expressed yesterday using the word ‘today’, he must replace this word with ‘yesterday’. —Frege, “The Thought”

1 Introduction

In “Demonstratives,” David Kaplan develops an account of the meaning of indexicals and sentences that contain them based on the concepts of content, character and context. The *content* of a statement is a proposition; which proposition a statement expresses depends not only on the linguistic meaning or *character* of the sentence used, but also on the *context*: who is speaking, to whom, when, where and in what circumstances.

In his essay, Kaplan briefly sketches an analogous concept of belief: we are in a belief state with a certain character in a certain context and thereby believe a certain proposition. As Kaplan says, we believe propositions *under* characters.

Kaplan then raises the following question. Suppose you have formed a belief of the sort that you would express with a sentence containing an indexical, say, “You are a computer scientist” or “Today is a nice day.” What do you need to do to *retain* such a belief, after you leave the context in which the sentence in question expresses it? What do you have to do to retain a belief that you once

*Forthcoming in *The European Journal of Analytical Philosophy*.

[†]Thanks to Pierre Jacob, Francois Recanati and Eros Corazza for comments on a much earlier version and to Elizabeth Macken for comments on the penultimate version.

expressed with “You are a computer scientist,” after the person you are talking to has left? What do you have to do to retain the belief you once expressed with “Today is a nice day,” after that day is gone?

Frege’s remark quoted above suggested a view to Kaplan: to retain a belief as one moves into a new context, one must adjust the character under which one holds the belief, to a new one that conforms to a sentence that would express the same proposition in the new context. There are patterns of character change that correspond to patterns of change in context of belief. To continue to think the same thing as the context of belief changes, is to think under a succession of characters that determine the same content in the succession of contexts, and conform to such a pattern. To think the thought I thought yesterday under the character of “Today is a nice day,” I must now think a thought under the character of “Yesterday is a nice day.”

Kaplan rejects this suggestion because of the case of Rip Van Winkle, who fell asleep for twenty years and woke up thinking he had slept for only a day. Kaplan thinks the Frege-inspired strategy would lead us to deny that when Van Winkle awoke, he had retained the belief he expressed, the day he fell asleep, with “Today is a nice day”. How would he express this belief? He might try to express it with “Yesterday was a nice day,” but this would fail. He would not have asserted anything about the day he fell asleep, but rather said something about the day before he spoke, a day that he slept through and of which he has no memory.

In his article “Understanding Demonstratives,” Gareth Evans found this an inadequate reason to abandon the Frege-inspired strategy. He seems to think that we should simply say that Rip has failed to retain his belief:

I see no more strangeness in the idea that a man who loses track of time cannot retain beliefs than in the idea that a man who loses track of an object cannot retain the beliefs about it with which he began ([?], pp. 87n-88n.)

In this essay I first review Kaplan’s theory of linguistic character, and then explain and motivate a concept of doxastic character. I then develop some concepts for dealing with the topic of belief retention and then, finally, discuss Rip Van Winkle. I come down on Kaplan’s side with respect to the Frege-inspired strategy, narrowly construed. But I advocate something like the Frege-inspired strategy, if it is construed more broadly. On my view it is remarkably easy to retain a belief, and I think Evans is quite wrong about Rip and Kaplan. The central concept I develop, however, that of an information game, is in the spirit of much of Evans’ work. I also borrow some of his terminology.

2 Linguistic Characters and Roles

The content of a statement is a proposition, “what is said” by the speaker in a literal sense. The context is a set of factors that determine what indexicals stand for: the speaker of the utterance, the time of utterance, the place of utterance, the circumstance or possible world in which it occurs. Characters are functions from contexts to contents. That is, a character takes a context as input (as its argument) and provides a content as output (as its value). These characters are mathematical representations of the rules that language assigns to expressions. Character is an interpretation of linguistic meaning.

Suppose that at a specific time in 1995—call it t —Kaplan says to Quine, “I live west of you now.” On Kaplan’s theory, here is what happens:

- The character of “I” is a function from a context to the speaker of the context. In this case, Kaplan is the speaker, and so he is the content of “I” in this context.¹
- The character of “you” is a function from a context to the person that

¹Kaplan develops his account within possible worlds semantics, and in that setting takes the value of an expression like “I” at a context to be a rigid *intension* rather than an individual. In this case, it would be an individual concept that picked out Kaplan in each world. This complication is basically an artifact of the semantical framework, not part of the intuitive set of ideas Kaplan tries to convey. I ignore this and other complications.

is addressed by the speaker of the context, at the time of the context, in the circumstances of the context. So the content of “you” in this case is Quine.

- The character of “now” is a function from a context to the time of the context. In this case, the content of “now” is t .
- The character of “live to the west of” is a function from a context to the 3-ary relation x lives to the west of y at t . This character is not sensitive to differences in context; its content is the same at all contexts.²
- The character of “I live to the west of you now” is built up out of the characters of the parts. It is a function that from a context to the proposition that a lives to the west of b at t , where a is the content of “I” in the context, b is the content of “you” in the context, and t the content of “now” in the context. In this case, the content is the proposition that Kaplan lives to the west of Quine at t .

Kaplan doesn’t assign characters exclusively to indexicals and sentences that contain them. Every expression gets a character. The characters of expressions that are not sensitive to changes of context are constant functions. The name “David Kaplan,” for example, is a function that returns David Kaplan at every context. I’ll call characters like this “loyal,” and characters, like that of “I” and “you” that are sensitive to changes in context “flighty”.

I like to interpret Kaplan’s characters in terms of what I call *utterance-relative* or *linguistic* roles. This interpretation links Kaplan’s ideas with an older tradition that emphasizes that indexicals are “token-reflexive.”³ Linguistic conventions assign a name like “David Kaplan,” to a particular individual. But they assign an indexical like “I” or “you” to an *utterance-relative* (or *linguistic*) role: being the speaker of u , being the being the addressee of u , etc. An

²I am ignoring tense.

³See, in particular, [?], [?] and [?].

utterance of an indexical does not stand for or refer to the role assigned to it; it stands for or refers to the object that plays that role, relative to the utterance itself.

Why are indexicals important? The answer has two parts. First, linguistic roles are closely associated with other roles that objects play in our lives. The speaker of the utterance I hear is often the person I am looking at; the place an utterance is made is usually the place the speaker occupies, and usually near the place the listener occupies. So, when I learn that an object plays a linguistic role relative to an utterance I hear, I learn about other roles that it plays.

Given this, indexicals are useful in two situations. Sometimes one wants to know more about an object that plays some linguistic role or associated role. I want to know more about the person I see before me. I ask: “Who are you?” Common sense and my facility at language assure me that the person I see before me will be the addressee of my utterance. So if I find out who the addressee is, I find out who the person before me is. I want to know the name of the city I find itself in. I ask: “What is the name of this city?” Common sense tells me that the city I am in is the city my utterance will take place in. So if I find out the name of the city in which my utterance occurs, I will find out the name of the city in which I find itself.

Sometimes I need to know what role some object is playing in my life, what its current relation to me is. I ask, “Who is David Kaplan?” You answer, pointing: “That man is David Kaplan.” I wanted to know which of the people I could see was David Kaplan. Common sense told you that I would be able to identify the person you demonstrated as one of the ones I saw. I ask, “When is July 4th?” You answer, “Tomorrow.” Common sense tells you that I will realize that your utterance occurs at the present time, and so that the day I am interested in is just 24 hours away.

When we hear an indexical, the first way of thinking of the referent that is afforded to us, is thinking of it as the thing that plays the utterance-relative role.

So when Quine hears Kaplan's utterance \mathbf{u} of "I live to the west of you now," the first way of thinking of Kaplan provided by this utterance is as the speaker of the utterance of "I". Quine's first grasp of the truth conditions of \mathbf{u} is something like, "The speaker of \mathbf{u} lives to the west of the addressee of \mathbf{u} at the time of \mathbf{u} ." This phase of understanding usually slips through one's mind without stopping, perhaps without rising to consciousness, as the utterance-relative roles give way to more interesting associated roles. In normal circumstances Quine will realize that the person he sees is the speaker of the utterance he hears, that that person is David Kaplan, that Kaplan is addressing him, and that the time of utterance is, for all practical purposes, the time the utterance is heard. So Quine will think something like, "Kaplan lives west of me now".

In other cases, however, the utterance-relative role may be our only way of thinking of the objects an utterance is about, at least until we have done some detective work. I find a faded note in my old copy of Wittgenstein's *Investigations*: "You are being an ass". I know that the note is true if and only if the person to whom it was addressed was being an ass at the time it was written. That may be the only grasp of the truth-conditions I can get, until I look through some papers and records and wrack my brains. Then I remember: someone passed the note to me, apparently thinking my remarks in a seminar on the private language argument were less profound than I did. Now I know that the note—or more accurately, the utterance for which the note was the token, was true if and only if *I* was being an ass.

Kaplan's system allows the possibility that we can say the same thing (utter statements with the same content) in quite different ways (using sentences with quite different characters). This happens if we use the sentences in contexts where, due to particular and perhaps peculiar circumstances, they turn out to have the same content. This happens when the same object plays two quite different utterance-relative roles. And it can happen without our knowing it.

Kaplan asks us to consider the sentences "My pants are on fire" and "His

pants are on fire”. The character of the first is a function that for a context with a as the speaker, returns the proposition that a ’s pants are on fire. The character of the second is a function that for a context with a as speaker, in a world in which a is pointing at b or otherwise calling attention to him, (“demonstrating b ”) returns the proposition that b ’s pants are on fire.

Suppose, to continue with Kaplan’s example, that he says, “My pants are on fire.” The content of his remark is the proposition that David Kaplan’s pants are on fire. Suppose now that he sees himself in a mirror, doesn’t realize that he is seeing himself, and, pointing at the man in the mirror, says “His pants are on fire”. The sentence is quite different, with a quite different character, but in this particular context the content is the same, that David Kaplan’s pants are on fire. Kaplan has said the same thing in two different ways without knowing it, although presumably he will recognize what is going on before long.

Now consider the difference between

- (1) My pants are on fire.
- (2) His pants are on fire.
- (3) David Kaplan’s pants are on fire on July 4, 1984 at 5 p.m.

One can think of (1) and (2) as tools. (1) is a tool for saying that oneself has burning pants. (2) is a tool for saying of someone that one can see and demonstrate, that they have burning pants. Since lots of people could, in principle at least, find themselves with burning pants, or find themselves in a position to demonstrate someone that has burning pants, these tools might be used, again and again, in different situations, to say different things about different people.

(3) doesn’t contain any indexicals, just names. Even the verb can be taken as tenseless, given the way the date is filled in. As we said, Kaplan assigns all expressions a character, not just indexicals, but for non-indexicals the character doesn’t do much. It’s just a way of making the theory work smoothly. So “Kaplan,” for example, is assigned a character that has David Kaplan as content

in all contexts. And (3) as a whole has a character that has the same proposition as content in all contexts, the proposition that Kaplan has burning pants on July 4, 1984. Thus the character of (3) is very loyal. It sticks with the same content in context after context. To put the point another way, with (3) we have a tool that allows us to express the same proposition, no matter when it is or where we are or to whom we are talking.

As a tool, (3) can seem a bit odd. Why would we need or want a tool for saying, just of one person, that he has burning pants? It seems like a very special purpose tool. We'll return to this question later.

3 Doxastic Characters

How can we conceive of beliefs, so that characters may be intelligibly assigned to them? One conception is that of belief as an attitude towards a sentence; the belief inheres the character of the sentence at which it is directed. I think this is a rather unsatisfactory conception. For one thing, it seems that fairly complex thoughts about things of all sorts would precede thoughts about sentences, both logically and psychologically.⁴

A more satisfactory way of conceiving of beliefs starts with the idea that they are concrete cognitive structures that arise in one's mind in certain situations; traces, as it were, of experiences of perception, learning, and inference. These structures have content; when one has a belief, there is (at least when things go right), something that one believes, a proposition P such that one believes P .

Beliefs so conceived will have two aspects that must mesh, *causal role* and

⁴I mention this conception mainly because many reasonably attentive readers seem to find it in "The Problem of the Essential Indexical" [?]. But it was not what I had in mind. See "Belief and Acceptance" [?]. I think this misinterpretation has three roots. I didn't properly distinguish between linguistic and cognitive roles. I underestimated the ease with which people would suppose that my view, that sentences with indexicals (or their characters), were more adequate ways of classifying belief states than were propositions, would have to rest on the view that sentences (or their characters) were what beliefs were directed at. Finally the problem referred to in the title had to do with the fact that indexicals *seemed* essential to *expressing* certain thoughts; from this some readers seem to have assumed that I thought that indexicals *were* necessary for *having* those thoughts.

content. The beliefs are caused by certain kinds of perceptions, and cause certain kinds of actions. But they also have a certain content; they are beliefs *that* such and such is the case, beliefs that P for some proposition P . These two aspects should mesh. A belief that P should cause action appropriate to its being the case that P , given one's desires and other beliefs.

By the *causal role* of a state I mean the various combinations of factors that bring the state about, and the various combinations of factors it brings about in turn. Consider the state of being nauseated. This state is universal and partial. To say it is universal is to say that different people at different times can be in the same state. You were nauseated last week, after a ride on the Ferris Wheel. I was nauseated yesterday, after binging on sushi and hamburgers. To say it is partial is to say that being nauseated constitutes only part of one's total state at a given time. I was nauseated and embarrassed and guilt-ridden and in a number of other states at the same time.

Each case of nausea has its own specific causes. But there are patterns. Certain combinations of factors bring about nausea, and nausea, together with other factors, has certain results. This pattern is the causal role of nausea.

On a given occasion, the causes of a specific case of nausea may be pretty similar to the causes of some other state. Perhaps the main cause of my nausea was pretty much the same thing as my feelings of guilt: massive overeating. But the effects of the two states are different. My nausea leads to me to take Alka-Seltzer; my guilt leads me to turn on the television. And in general, the patterns are not the same. Lots of things cause nausea that don't cause guilt and vice versa, and nausea, in combination with various factors, causes lots of things that guilt, in combination with those same factors, wouldn't cause.

One can think of causal roles in various strict and philosophical ways or in a sort of loose and casual way. The second will suffice for my purposes. Think of the causal role of a state as its typical causes, the things you would expect might cause an instance of that state, in more or less ordinary circumstances,

and similarly with its typical effects.

We said that the causal role of a belief should mesh with its content. But Kaplan's case shows that this meshing is not simple and direct.

First, note that the causal roles of the states we imagine to occur are quite different. If we heard that someone believed that their very own pants were on fire, we would expect that something like this happened. First, they dropped an ash from a pipe into their own lap, or stood too close to a fireplace, or something like that. Then they felt some unaccustomed warmth in their nether regions. Then perhaps they smelled something like wool or cotton or rayon burning. Then they looked down where things were heating up, and saw the smoke. Those are typical causes of the state one is in when one says, "My pants are on fire." The typical effects would include strong emotions like fear, and attempts to douse oneself with water or put out the fire in some other way, and saying, "My pants are on fire."

Now consider the state Kaplan was in when he saw the man in the mirror, who just happened to be him, with his pants on fire. This state is typically caused by seeing smoke and flames erupt from the trousers of someone standing in front of one. It typically leads to concern and the attempt to help and shouting, "His pants are on fire". So the typical causes of this state are quite different from that of the first state.

Next, note that these different belief states do not line up directly with propositions believed, so there is not a simple and direct meshing between causal role and content. We can imagine a lot of people being in the "My pants are on fire" belief state; they wouldn't thereby believe the same propositions. Kaplan would believe that Kaplan's pants were on fire, Searle would believe that Searle's pants were on fire, and so forth. And we can imagine a lot of people being in the "His pants are on fire" state. They all have something in common, but it's not what they believe—it's not the content. They would believe quite different propositions, depending on whom they were looking at.

At this point, it is very helpful to postulate a level of meaning for beliefs that is *analogous to* but not *derived from* that of character for sentences—characters that in fact are quite independent of language. A belief will have a certain content because (i) it has a certain character, and (ii) it occurs in a certain belief context—it is a belief held by a certain person, at a certain time, in a certain place, attending to certain objects, etc. The virtue of thinking of a belief in this way, is that it makes intelligible how this meshing between causal role and content could take place, in a systematic psychology of content, for these characters will correspond more closely to belief states than propositions do.

A wider class of roles is needed to characterize beliefs than is needed to characterize utterances. We noted that the characters that are associated with indexicals are based on *utterance-relative roles*. These roles are based on relations objects might have to a given utterance. To characterize beliefs we need what we might call *thought-relative* or *cognitive* roles. These roles are based on relations that an object can have to a given episode of thought or a particular belief, such as being the owner of the thought, (the self role), being attended to by the owner of it, being remembered by the owner of it, being held in the right hand of the owner of it, being above the owner of it, being sat on by the owner of it, and so forth.

I'll indicate doxastic characters by quoting the sentence a speaker might use to communicate the beliefs one has under the character. Of course the characters of these sentences are linguistic rather than doxastic; I supplement the indexical expressions with bracketed material to identify the underlying cognitive role involved. For example, “that man [the one I see]” and “that man [the one I remember]”, “I [self]” and “now [the moment of thought]”.

There is a difference between referring to someone as “you” and designating them with the phrase “the person I am addressing,” even though the role of being the addressee is involved in both cases. When one refers to a person as “you”, one does not *say* that one is addressing them; one exploits that fact to

refer to them. One can learn how to use “you,” without being able to explicitly formulate the conditions under which a use of that word refers to a person (perhaps because one has not thought about it, or has not yet attained the concepts, such as the concept of a word and the concept of reference, that would be necessary to think about it). Even so, one has *some* grasp of these conditions; one has some sort of positive doxastic attitude towards the fact that a use of “you” refers to the addressee. But one may not be properly said to *believe* that the reference of “you” is the addressee. One knows *how* to use “you” to refer to the addressee, even though one may not know *that* a use of “you” refers to the addressee. In these cases, I say that one is *attuned to* the way “you” works.

Similarly there is a difference between being able to think of a thing or person in virtue of some role it plays in one’s life, and being able to articulate that role in thought or speech and think of it *as* the thing or person playing that role in one’s life. Consider a child, who is thinking about a dog that she saw an hour or so before. She has a certain memory of the dog, and it is in virtue of this memory that she is able to think of the dog. I would represent the character of her thought with “That dog [the one I remember] was very cute”. This is not the same as thinking “The dog I remember was very cute”. The child might lack the concepts needed to think this thought; even if she has them, it may take a bit of time and wit to figure out that the dog can be characterized as the one she remembers. Still, the child would be attuned to the fact that the dog she is thinking of is one that she remembers, in that she would know how to consult her memory for more information about the dog: to find the dog she would go back to where she had last seen it, for example.

Now let’s return once again to Kaplan and his burning pants. We imagined Kaplan making two sets of observations of his own burning trousers. We imagined him using the sentences, “My pants are on fire” and “His pants are on fire,” to express what he believed. But he might make the observations,

and acquire the beliefs, without saying anything. Even if they don't lead to distinct utterances, the belief states are quite different. If Kaplan notices that his pants are on fire in the usual way, he will be in quite a different state than if he notices a man with burning pants in a mirror. The difference in the actions he would take in these circumstances, including the difference in sentences he would utter if he were to put what he noticed into words, *reflects* a difference in beliefs. The concept of believing under a character is intended to capture this difference. The difference in belief would be there, even if Kaplan didn't say anything. And the difference between the two cases is not the proposition that is believed, but the character under which it is believed.

There are two characters here, one corresponding to that of "My[self] pants are on fire" and one corresponding to that of "His[the man I am looking at] pants are on fire". These characters come much closer to lining up with causal roles than do the propositions believed.

Thus two things stand out at the level of character that get obliterated at the level of content. First, the difference in causal role of the two beliefs. Second, the common nature that different beliefs with different contents belonging to different people at different times might have, and in virtue of which these different situations might instantiate the same psychology of content.

Recall the distinction between flighty and loyal characters from our discussion of linguistic characters. Loyal characters yield the same content from any context; these characters are the sort that belong to sentences that don't contain indexicals. Flighty characters yield different contents in different contexts; these characters are the sort that belong to sentences that contain indexicals.

The same distinction applies to doxastic characters. "David Kaplan has burning pants at 4 p.m. on July 4, 1996" is a very loyal character; "I [self] have burning pants now [the present moment]" is a very flighty character, expressing a different proposition with each variation in time or thinker.

If we know that an utterance of sentence with a flighty character expresses

a certain proposition, we may know quite a bit about the context in which the utterance occurred. If an utterance of “I am a logician” has the content that David Kaplan is a logician, then the speaker must be David Kaplan. If an utterance of “That man works for Microsoft now” has the content that Elwood Fritchey works for Microsoft on October 27, 1995, then the utterance occurred on October 27 and the speaker was someone in a position to demonstrate Elwood Fritchey. The objects that get referred to by the indexicals must be playing certain linguistic roles, and hence certain roles in the life of the speaker.

The same is true with beliefs. If I believe that David Kaplan has burning pants under the character “that man [the one I see] has burning pants,” then I must be looking at David Kaplan. If anyone believes that proposition under the character “I [self] have burning pants,” that person must be David Kaplan.

If my beliefs are under characters, it seems that many of the characters they are under are very loyal and not very flighty. Right now I believe that David Kaplan is a logician. I am not speaking to him, seeing him, or hearing him right now. He is playing no role in my life that would enable me to refer him by uttering “you” or by demonstrating him and uttering “he.” What character do I believe these things under, and what is the point of such beliefs? What causal role, if any, do they have? Why do people have such beliefs, and why do others care whether or not they do?

4 Information Games

We ended each of the last two sections by wondering about loyal characters. Why do we want sentences that have loyal characters—what use do we have for such special purpose tools? And what sorts of beliefs have loyal characters, and why are they important? To consider these questions, I want to introduce the concept of an information game.

An information game involves the acquisition and later application of a belief about an individual. That is, at some time one comes to believe something about

some person or object. Then, later, that belief guides one's behavior towards that object or at least an object that one takes to be the same as it. I call the object about which one acquires the belief the *source*.⁵ I call the object to which one applies the belief—the object one takes to be the source—the *applicandum*. In any information game, one faces the problem of making reasonably sure that the source is the applicandum.

Suppose I meet Elwood Fritchey at a party; he tells me he is a programmer for Microsoft and I believe him. So I acquire a belief. Later I ask him if he knows why the Macintosh version of Word 6.0 works so slowly. I direct this question to him because I believe him to be a Microsoft programmer. If I didn't believe that, I wouldn't ask him. It doesn't make much sense to ask some random person this question. My belief that he is a Microsoft programmer is part of the reason I ask him; my asking him is an application of the belief.

In an information game I acquire a belief, and then later I apply the belief I acquired to a certain object: because I have the belief, and take the object to be the one the belief is about, I deal with it in a certain way. In this case, I ask a person a certain question, because I think he is in a position to know the answer.

I will describe eight information games, which I call “straight-through”, “tracking” “detach-and-recognize”, “updating”, “recollection”, “inference”, “planning” and “communication”.

I introduce the first three by considering three scenarios for what happens in between my acquiring the belief and my applying it.

First, virtually nothing comes between. I am talking to Fritchey, face to face. He says to me, “I am a programmer at Microsoft.” I acquire a belief I would express with “You are a programmer at Microsoft.” I say, on the basis of that belief, “Why does the Macintosh version of Word 6.0 run so slowly?” This is the straight-through information game.

⁵See [?] for the term “source” used in more or less this way.

I stand in a certain relationship or family of relationships to Fritchey. He is in front of me; I am looking at him; I am talking to him; he is talking to me. When someone has these relations to me, or as I shall put it, occupies these *agent-relative roles* in my life, I have ways of finding out information about him. (I'll say that there are epistemic techniques associated with the role, or for short, that it is an epistemic role.) I can look, and I can ask and listen to the answer, to mention the most obvious ways of gathering information about the person one is talking to face-to-face. There are also ways of acting, the success of which will turn on the characteristics of the person that plays this role. (I'll say that there are pragmatic techniques associated with the role, or for short that it is a pragmatic role.) The action of asking, "Why does Word 6.0 work so slowly," with the goal of finding out the answer, has a chance of being successful if the person I am talking to is a programmer at Microsoft.

In the straight-through information game, making sure that the source is the applicandum is not a big problem. The epistemic role and the pragmatic role are closely associated and the time difference between acquisition and application is very short. There is not time for one object to cease to play it and another object to take its place.

In particular, I don't need to know a lot about Fritchey to successfully play the straight-through information game. I don't need to know his name, or have any way of recognizing him. I just need to be able to tell that only one object plays the role in question in my life for the few moments it takes for the game to be played.

Second case. I break off my conversation with Fritchey, after learning that he is a Microsoft programmer but before asking him about the Macintosh version of Word 6.0 . But I keep my eye on him as he and I move our various ways around the party. Other people assume the role in my life that Fritchey had: that is, I converse with them. Fritchey plays a succession of roles: object to my left, object to my right, object I am glancing at. Later we end up face to face

and I ask him, “Why does Word 6.0 run so slowly”.

Here the task of making sure that the applicandum is the source is more complicated. I have to *track* Fritchey, to use another apt term of Evans’s ([?]). This requires that I be able to ascertain that a single object has played a succession of roles in my life. This requires more than is required for the straight-through language game, but it doesn’t require that I have the ability to recognize or re-identify Fritchey if I lose track of him.

In the straight-through and tracking games, the player stays *attached* to the source, in the sense that he remains in an epistemic relation to the source from acquisition to application. In the straight-through game it is the same relation and virtually the same time. In the tracking game it is a succession of relations through an interval of time.

The third game I call the “detach-and-recognize” game. I learn that Fritchey is a programmer for Microsoft. Then I go home. I don’t see him for days or even years. At some later point I see him again. I recognize him, remember his job, and ask him, “Why does the Macintosh version of Word 6.0 run so slowly?”

After I leave the party, Fritchey is not playing an epistemic or pragmatic role in my life. He is of course still related to me. He is a certain distance and direction from me, but I am not aware of it; it is not part of my conception of him. There are various actions I could take to find out things about him, or to have an effect on him. But there is no simple technique, like looking or asking a question that I can use to find out more about him, and no simple technique like talking to him or shoving him or gesturing to him that I can use to have an effect on him.

Nevertheless, Fritchey still does play a role in my life. He is the man of whom a certain thought of mine is a memory. This role, this relation to my thought and to me, gives me a way of thinking of him. It is a very loyal character. I can carry the memory with me as I travel away from Fritchey and he from me; its being a thought of Fritchey depends on the way I acquired it, but not on my

present circumstances.

As an analogy, consider the note I left in my copy of the *Investigations*. I don't remember who wrote it. The signature is hard to make out. The writing, never very legible, is faded. In spite of all this, the person who played the role, "writer of this letter," is the same as it was thirty years ago. It is a very loyal role. And the source of memories is similarly loyal.

In such a case, what is the point of continuing to believe anything about Fritchey? It doesn't have much of a point, unless there is a good chance I will encounter him again and be able to apply the belief usefully then. But of course this is something we do a lot. We acquire a belief about an object at one time, when the conditions are favorable for doing so. We apply the belief at another time, under quite different conditions. These later conditions might make it impossible to acquire the information if we did not already have it.

Doing this requires that we be able to recognize the object in question. Unless I remember a little bit about Fritchey, so I can recognize him the next time I see him, I won't be in a position to use the information I have retained. The information that we need, to make detach-and-recognize a viable information game, falls into two categories. There is the information we will use when we encounter the object again, to decide what to do. The information that Fritchey is a programmer for Microsoft is what leads me to ask about the sorry state of Word 6.0. Call this sort of information, "information for action." But there is also "information for recognition." My tidbit of information about Fritchey will sit useless in my brain, unless I remember also that he is a big man with a red beard, etc., information that will allow me to recognize him next time I meet him. Of course, these are different functions or uses for information, not different categories of information itself. In many cases, a bit of information will serve both purposes.

To be useful, then, a detached belief, like the belief that Fritchey is a programmer for Microsoft, will have to be part of a larger *file* on Fritchey. The

additional elements in the file provide the facts about him that might enable me to recognize him.

Nevertheless, there is no guarantee that beliefs will be useful. Consider my belief that Fritchey was a microsoft programmer. Under what character do I continue to believe this, once I detach from Fritchey and leave the party? Simply the character:

That man [the source of the belief] was a programmer for Microsoft then [the time the belief was required].

The belief is similar to the note I discovered in my old copy of Wittgenstein's *Investigations*. By itself, without connection to and augmentation from other sources, it is useless. This belief, by itself, includes no information about Fritchey's name, for example, so I couldn't even use it to contact him if I had a directory of Microsoft employees.

A belief that is detached and virtually useless can nevertheless be quite conscious, even vivid. The frustration that such beliefs can engender is well-known, and we have linguistic devices for expressing them. The following conversation not only makes perfectly good sense, it is all too familiar for some of us:

“Do you know any programmers for Microsoft?”

“Yes...uh, uh, what's-his-name was a programmer for Microsoft. I don't remember his name.”

“When did you meet him?”

“I'm not sure.”

“Where did you meet him.”

“I don't remember that either. I can't remember much of anything about him, except that he was a programer for Microsoft.”

“Thanks.”

A belief like this one, totally useless at one time, may become useful later. Recovered memories, or outside sources, may disclose more information about the person I remember. Then the apparently useless bit of information may prove quite useful.

The detach-and-recognize game provides part of the answer to our question about the point of beliefs with loyal characters. Such beliefs serve as components of larger, useful, beliefs, when the objects the beliefs are about are recognized. They are completions for pragmatically attached beliefs, broken off from earlier epistemically attached beliefs.^{6,7}

Detached beliefs do not simply sit in our minds and gather dust while we wait to encounter the objects they are about. They are parts of various activities that do not directly affect the objects they are about, but may have profound indirect effects on them. We use such beliefs in five additional information games.

In the *recollection game*, I try to squeeze more information out of my memories. I may have formulated the explicit belief that Fritchey was a programmer for Microsoft as I talked to him; he leaves and I detach the belief, thinking of him as “that man [that I remember]”. But associated with the belief, at least for a while, are memory images from which I can cull some more beliefs, that may aid in recognizing him—e.g., that he had a red beard and wore green suspenders over a plaid shirt and had an unusual fascination with tupperware.

I do not need to be attached to Fritchey to recollect more about him, and I also do not need to be attached to him to play the *inference* game—to draw inferences on the basis of my beliefs about him. I combine these beliefs with other beliefs I have, to flesh out my conception of Fritchey. Given his job, I

⁶Can I find the burge article to cite here?

⁷What we have described so far is perhaps the main variation on the detach-and-recognize information game. There is a source of my belief, or of my file of beliefs, about a certain person. But we can also have sourceless files. Sometimes we can figure out quite a bit about an object that meets a certain description, and work up a “file” about it, even though we have never encountered it or communicated with anyone that has. In this case, once we have done the reasoning, we are left in a situation analogous to having examined a source.

infer that he is bright and likes computers. Given his dress, I assume that he is happy in Seattle and enjoys being outdoors. Given that he worked for Microsoft when I talked to him, I infer that he will probably continue to do so for some length of time.

I call a rather special kind of inference *updating*. This is an inference made not on the basis of observed or inferred movements or changes in the thing my belief is about, but on the basis of changes in my own situation, or general changes, like the passage of time. When I meet Fritchey, I think “This man [the man I see] is a programmer at Microsoft now [the time of the thought].” Later I think “That man [the man I remember] was a programmer at Microsoft then [the time I remember].” The change from “now [the time of the thought]” to “then [the time of the thought]” does not represent an inference about how Fritchey has changed since I last saw him. I am just updating; changing the character in a way that preserves the content, given the new circumstance of belief.

This update is quite different from another additional thought I may have. Given the relative permanance of jobs at Microsoft, I may figure at the later time “That man [the man I remember] is a programmer at Microsoft now [the moment of thought].” This is not an update that preserves the previously believed content, but an inference—although perhaps one that is not very risky, if the time interval is short.

It is possible that I have met Fritchey before. Or perhaps I have read an article about him—let’s suppose he has been involved in the Seattle Commons project and was pictured and briefly described in an article about it. I already have some concept of him. Upon reflection, I may find that this concept fits the man I thought I had just met for the first time too nicely to be coincidental, and conclude that Fritchey is the man I read about.

When I read about Fritchey in the paper, I was playing the communication information game. Someone else had met him and talked to him and taken his

picture. They had been in an epistemic relation to him, and acquired information about him. The information I get is detached.

When I read about Fritchey, he plays a special sort of epistemic role in my life, one mediated by symbols. He was the man I was reading about; the man referred to by the words I saw in the paper. When I quit reading the paper, I continued to have beliefs about him—detached beliefs. For a while, my beliefs about Fritchey may have been tied to memories of the article or the paper. But often we retain beliefs about a person or things we have encountered or read about, long after the memories of the perceptions, conversations or text from which we learned about them have faded to insignificance. The role these objects play in our lives is just that of being the source of the beliefs we have about them.

The beliefs I gained about Fritchey from reading the article may be quite detached and fragmentary. Perhaps I don't remember his name; I just have a concept of a man who was a programmer at Microsoft, worked on the Commons, wore green suspenders, and had a bit of an odd name that reminded me of a British Rock Star of an earlier era. Such beliefs would seem quite useless. But there could be enough there to make a pretty plausible conjecture, going over things after meeting Fritchey at the party, that it was the same man. And then I would have a belief that, next time I met Fritchey, could lead to a question about his work on the Commons, as well as the one about Word 6.0. The fragmentary beliefs turn out to be useful after all.

So now I am sitting at home. I acquired and detached a belief about Fritchey. I have recollected and inferred. And now I intend to ask him, next time I see him, about Word 6.0 and Seattle Commons. I am not now in a position to do anything to him directly; he plays no pragmatic role in my life at the moment. But he will, and I am forming plans about what to do then—playing the planning information game.

All of these last five games, recollection, inference, updating, communication

and planning, could be played while I was still attached, epistemically or pragmatically, to Fritchey. But our main interest in them is that they can be played with detached beliefs, beliefs with loyal characters, and help us see the point of having such beliefs. Beliefs with loyal characters then have a causal role that has to do with the kinds of inferences and plans they lead to, and their effects on action when combined with other attached beliefs.⁸

5 Misidentification

Suppose that a few days later I mistakenly take someone else to be Elwood Fritchey— Elwood’s brother Alphonse, say. I see Alphonse, as tall and red-bearded as Elwood is, one day at the feed lot. “Elwood,” I say, talking directly to Alphonse, “Why does the Macintosh version of Word 6.0 run so slowly?”. Here I have not only lost track of Elwood, I have applied my Elwood belief to someone else.

This is no reason to say that I have lost the belief. My belief, the one I acquired at the party, that Elwood is a Microsoft programmer, is part of the explanation for my question to Alphonse—it is a misapplication of that very belief. Since the belief explains the mistake, the mistake can hardly be reason for supposing the belief to have been lost. I asked Alphonse the question, *because* I believed that Elwood was a programmer at Microsoft, and I took Alphonse to be Elwood.

Suppose I say, by way of explaining my question to the puzzled looking Alphonse: “You are Elwood Fritchey, the Microsoft Programmer”. This would be an indirect speech act, asking him if indeed he was Elwood. But consider the literal statement I make. Clearly it is false. The fact that I intend to be speaking to Elwood and expressing a belief about Elwood does not change the fact that

⁸I am here concerned with what seems to be a rather basic part of our conceptual scheme, beliefs about individual things and people. Perhaps as an explicit disclaimer I should note that there are other important classes of beliefs to which such detached beliefs are relevant. For example, a detached belief can serve as a confirming instance or a counterexample to a generalization.

the person who is playing the relevant linguistic role in my life is Alphonse.

Do I have, at this point, a belief in the singular proposition, that Alphonse Fritchey is a Microsoft Programmer or the proposition that Elwood is Alphonse? I have said these things, but do I believe them? It is not necessary to say that I do. These beliefs are not required to explain my question; it is adequately explained by my true belief about Elwood Fritchey and my false belief that I am talking to Elwood Fritchey. Those beliefs explain why I think I can say something to and about Elwood Fritchey by using a sentence that contains “you”.

There are reasons for saying that I don’t believe these propositions. When the mistake is discovered, I would quite naturally say, “I thought you were Elwood” or “I took you to be Elwood”. It seems that not every positive doxastic attitude is comfortably described as a belief. My mistake, in taking Alphonse to be Elwood, may *lead to* false beliefs, if it remains undiscovered. But if it is simply a transitory thought, my true belief about Elwood and my false belief that this man was Elwood suffice to explain my false statement.

6 Thinking about days

Detach-and-recognize is a reasonable strategy for dealing with the fact that individuals come and go. They become indexically inaccessible, cease to play any epistemic or pragmatic role in our life for a while, and then re-enter it again. In the meantime, by detachment, we can retain our beliefs, update them, use them as the basis of inferences, make plans about future interactions and communicate with others.

Keeping track of the relative locations that times and immobile things have to us is easier than keeping track of people and other mobile things, at least in principle. If I keep track of my own movements in space, then I can correct for those movements and still think of all the places I have been and buildings and landmarks I have seen in terms of their relative distance and direction from me. But I don’t usually do this. Right now, for example, I am working at my

home. I remember walking at Palo Alto's Baylands a few hours ago. I don't think of the Baylands in terms of their relative distance and direction from me. I couldn't say right off the top of my head whether they were to my left or right, in front or behind me. If I want to think of them as "over there" I would stop and consider where I live in relation to them, and what my orientation in my house is, and figure it out from there. I might do this for example if I saw some fireworks, and wondered if they could be coming from the baylands—if the baylands were over there [where I am looking].

We do a much better job keeping track of whether each of our experiences lies in the past, present or future. This is fairly easy because there is an exact correspondence between the mode of thinking about the experience (remembering it or planning to have it or having it) and its position in time relative to the present moment. But we don't usually keep very precise track of events in terms of these cognitive roles. If I have an important appointment coming up, I will definitely be aware that it is in the future and not the past; I will probably remember the time. As the day goes on I may occasionally figure out how long it is before the event. But I don't retain my belief by constantly updating in terms of "minutes from now". I remember that the appointment is at 5; I look at the clock and see that it is 3:30, and then I think, "An hour and a half from now I see the dean" or something like that.

We can think of "now [the present time]" and "today [the present day]" as ways of thinking about days that are both pragmatically and epistemically attached. On September 1st I can find out what September 1st is like in Palo Alto by looking. On September 1st I can make September 1st a day on which I take a walk by taking a walk.⁹

On September 2nd, I can no longer make any changes in what happens on September 1st.¹⁰ I am not pragmatically attached. I am no longer in a position

⁹See [?]

¹⁰Setting aside "Cambridge changes." I can talk about September 1st on September 2nd, thus making it, what it hadn't been before, a day talked about by me on September 2nd. But this doesn't change what September 1st was like; it doesn't change what happened on that

to actively explore September 1st. I am not epistemically attached. I can, however, still explore my own memories, and may form new beliefs as a result of that exploration.

How about tomorrow? I am not pragmatically attached to tomorrow. I can't do anything today that will directly change what happens tomorrow. I can indirectly affect tomorrow by making changes today in myself and others that will have an effect tomorrow. I can plan a trip to a concert, or call a meeting. And I am not now perceiving tomorrow's events; I am not epistemically attached. If I am not now, and never have been, epistemically or pragmatically attached to tomorrow, how can I think about tomorrow, reason about it, make plans for it, and the like?

Our system of dates provides detached ways of thinking about days that we have not yet encountered. This system exploits important metaphysical differences between days on the one hand and things and people on the other. We can systematically talk about days we have not encountered and keep track of them in our thoughts and plans, simply in virtue of their position in the sequence of days, as reflected by the calendar. This is in part because days are connected to one another in a linear, predictable, fashion, dependent only on the most general facts that circumscribe existence on earth. It is in part because of the somewhat puzzling fact that the date of a day, its position in the calendar, although seemingly a relational fact about it, seems also to be its most essential feature. We can hold a day in our modal imagination and change everything about what happens on that day. We can suppose that all sorts of things will happen on July 4, 1997; we can coherently imagine July 4, 1767 having been quite different than it is. We can even suppose that everything that will happen on July 4, 1997 could have happened instead on July 4, 1767, due to the absence of some sequence of events a couple of million years ago that delayed everything 230 years. (In this case, of course, July 4, 1767 would be called "July 4, 1997".)
day.

But we cannot consider a given day as being in a different place in the order of days than it actually is. At least it is not easy.

Can next July 4th be a *source* of my thoughts? No, for it lies in the future, and cannot be part of the cause of my beliefs. But I can have a *sourceless* detached belief about next July 4th. The combination of our system of dating, our holidays and our traditions allows me to figure out quite a bit about next July 4th, even though I have not causally interacted with it.¹¹

Consider what is involved in a simple matter of arranging a meeting. We decide to meet on a certain day—say two weeks hence. The decision is recorded on a calendar. At this point I am not in a pragmatic relation to the day. I have made a plan for what will happen on the day, but I cannot make it happen. I can't attend or not attend the meeting until the day itself arrives.

The task of recognizing a day is disanalogous in many ways to recognizing an individual. When we can identify today, yesterday, or tomorrow by its date, it is because we have been keeping track, or are exploiting someone else (like the newspaper publisher), who does.

Notice here that “keeping track” does not mean following a particular day. To keep track of time is usually not to pick out a day and track it, but to be aware of which day it is. That is, to be aware of the important properties of the day that plays the “today” role—what the date is, what events are planned, etc. To lose track of what day it is, means not knowing that the day that plays that role has some other important attributes, like being one's anniversary or the day a philosophy paper was promised to an editor.

For the most part, apart from human activities and human institutions, like calendars and the dates on newspapers, days don't have particular features, easily ascertainable by the ordinary citizen, that set them off from other days, in roughly the same season. When Rip awakes in the hills, there is nothing about the look of things that suggests that it is a 1786 day rather than a 1766

¹¹I do not deny that we can and do have sourceless beliefs about persons and things as well as times and places.

day. It's only when he gets to town, and observes humans and human artifacts, that things begin to not fit.

So now let us review the analogies and disanalogies we have found in terms of our information games. There are clear analogues, with days, to the straight-through and tracking games. There is not a direct analogue to the detach-and-recognize game, since we cannot re-encounter a day once time has left it behind. However, the peculiar metaphysical status of days, that is reflected in our system of dates, and permits us to arrive at sourceless beliefs about them, provides an alternative route to detached thinking detached thinking, planning, information exchange, and the like.

With this all in mind, let us, finally, return to Rip Van Winkle.

7 Rip Van Winkle

What should we say about Rip Van Winkle?

Rip Van Winkle acquired a belief the day he fell asleep—July 3, 1766, say—a belief that that day was a fine day. He held this belief under the character “Today [the day of this thought] is nice.” Then he slept for twenty years and two days, until July 5, 1786, and walked back to town.¹² What happened next?

The possibility that struck Kaplan and Evans is that Rip merely updated his belief. On July 3rd he never forms any explicit belief other than “Today [the day of this thought] is a nice day”. When he awakes on July 5th, the belief is updated, due to his awareness of having slept through a night, and his lack of awareness of having slept twenty additional years, to “Yesterday [the day before the day of this thought] was a nice day.” He falls out of epistemic contact with the current day when he falls asleep, but has a ready-made character in mind

¹²Actually, Knickerbocker does not tell us the exact dates, just that the Declaration of Independence and the American Revolution occurred while Rip slept. I've added a day to his twenty year sleep for expository convenience; the important thing is that Rip fell asleep on July 3 and awoke on July 5 thinking it was July 4; aside from the possibility of connecting humble analytical philosophy with important literature, the extra twenty years doesn't matter much.

for when he wakes up. But then what is there left of the original belief except the false one about July 4th? But the false belief can not be the true belief, so hasn't Rip lost the belief in question? This seems to be the argument that threatened Kaplan and appealed to Evans.

But even in the case of such thin updating, there are backup characters for Rip to hold his belief under. When Rip believes, towards evening, as it grows dark, "Today [the day of this thought] was a nice day," he has memories of seeing the flowers and feeling the sun, and so forth. So the character, "That day [the day I remember] is or was a nice day" is available to sustain his belief, when the attempt at updating goes awry. Even if these memories fade, there is the character, "That day [the day this belief was acquired] is or was a nice day."

So my view is this. When he awakes on July 5th, Rip updates his belief according to his view of how the context has changed. His view about the change of context is mistaken, and the new character, "Yesterday [the day before the day of this thought] was nice" is not a way of believing the original content. But that is no reason to say that Rip has lost his original belief. He retains it under various other backup characters.

That's what I have to say about Rip; what about the other concrete characters that figure in our story, Frege, Kaplan and Evans?

We can't pin anything much on Frege, for we can't hold him responsible for the strategy about belief that his remark on saying inspired. We could consider whether he was right about saying, but we won't do that in this essay.

The strategy that Frege's remark suggested to Kaplan and Evans is that retaining belief consists in moving from flighty character to flighty character in ways that reflect change in context. I think I have refuted that strategy.

But my own strategy is the broad interpretation of the Frege-inspired one, generalized and freed from its association with his particular example. The detach-and-recognize strategy for handling information itself embodies a regu-

lar transition, from strong characters to loyal characters and, when recognition occurs, back to strong characters. This Frege-inspired doctrine, like most, perhaps, is inadequate when construed narrowly, plausible when given a broader interpretation.

We can't fault Kaplan for thinking that there was more to be said about Rip Van Winkle for in this he was correct. Evans was wrong, I think, about Rip and about Kaplan. But much of what I have said is similar in spirit to ideas one finds in the body of his work. I will end with a couple of remarks on interpreting Evans.

In approaching Evans, it seems to me one must try to separate his own information-oriented approach to things from the devotion to a version of Frege filtered through Davidson that crops up now and then, most especially in "Understanding Demonstratives," the essay from which the quote above was taken.

Davidson's reliance on Tarski and T-sentences in explaining his views on meaning has inspired a tradition in semantics that one might call "homogeneous meaning explanation." We explain the meaning of a sentence by using the same sentence, or one with the same meaning. The work of Kaplan belongs squarely in the tradition of "heterogeneous meaning explanation". One explains the meaning of a sentence by showing, using whatever language one might want, what sort of tool it is, how it conveys various things in various circumstances.

Imagine that you are explaining a Mercator projection map to a child. One way of explaining, perhaps, is to use another Mercator projection map. A better way is to use a globe. I think the benefits of heterogeneous over homogeneous explanation are similar.

As Davidson himself notes, the homogeneous strategy works with indexicals only when one supposes that the semanticist's explanation and the statement explained are made by the same person at the same time.¹³ One interpretation of "Understanding Demonstratives," is that Evans tries, using a variety of ideas

¹³See [?], last section.

that are interesting in their own right but are not necessarily well-suited to the purpose, to extend the homogeneous treatment of indexicals beyond the special case in which it works. I think this experiment fails. And I also think that the information-oriented approach that one finds in much of Evans work, and in parts of “Understanding Demonstratives,” is basically heterogeneous in its implications for semantics.

References

- [Almog, 1989] Almog, Joseph, John Perry and Howard Wettstein, eds. *Themes From Kaplan*. New York: Oxford University Press, 1989.
- [Barwise and Perry, 1983] Barwise, Jon and John Perry. *Situations and Attitudes*. Cambridge: MIT-Bradford, 1983.
- [Burks, 1949] Burks, Arthur. Icon, Index and Symbol. *Philosophical and Phenomenological Research*, vol. IX, no. 4, June 1949: 673-689.
- [Crimmins, 1992] Crimmins, Mark. *Talk About Beliefs*. Cambridge, Mass.: Bradford Books, 1992.
- [Evans, 1981] Evans, Gareth. Understanding Demonstratives. In Herman Parret and Jacques Bouveresse (ed.) *Meaning and Understanding*. Berlin and New York: Walter de Gruyter, 1981: 280-303. Reprinted in [?], pp. 71-96.
- [Evans, 1973] Evans, Gareth. The Causal Theory of Names. *Aristotelian Society, Supplementary Volume* 47 (1973), pp. 187-208.

- [French, et. al., 1979] French, Peter A., Theodore E. Uehuling, Jr., and Howard K. Wettstein, editors, *Contemporary Perspectives in the Philosophy of Language*. Minneapolis: University of Minnesota Press, 1979.
- [Israel, Perry and Tutiya, 1993] Israel, David, John Perry and Syun Tutiya. Executions, Motivations and Accomplishments. *The Philosophical Review* (October, 1993): 515-40.
- [Kaplan, 1989b] Kaplan, David. Afterthoughts. In [?]: 565-614.
- [Kaplan, 1989a] Kaplan, David. Demonstratives. In [?]: 481-563.
- [Kaplan, 1978] Kaplan, David. Dthat. In [?], pp. 383-400. Reprinted in [?]: 11-33.
- [Kaplan, 1979] Kaplan, David, On the logic of demonstratives. *The Journal of Philosophical Logic*, 8 (1979): 81-98. Reprinted in [?]: 401-412.
- [Perry, 1994] Perry, John. *Davidson's Sentences and Wittgenstein's Builders*. APA Proceedings, 1994.
- [Perry, 1993] Perry, John. *The Problem of the Essential Indexical and Other Essays*. New York: Oxford University Press, 1993.
- [Perry, forthcoming] Perry, John. Indexicals and Demonstratives. In Bob Hale and Crispin Wright, eds., *Companion to the Philosophy of Language*, Blackwells, forthcoming.

- [Perry, 1980] Perry, John. A Problem about Continued Belief. *Pacific Philosophical Quarterly* 61, no. 4 (1980): 317–22. Reprinted in [?]: 69-90.
- [Perry, 1979] Perry, John. The Problem of the Essential Indexical, *Nouûs* 13, no. 1 (1979): 3–21. Reprinted in [?]: 3-49.
- [Reichenbach, 1947] Reichenbach, Hans. §50. Token-reflexive words. In *Elements of Symbolic Logic*. New York: The Free Press, 1947. Pp. 284ff.
- [Yourgrau, 1990] Yourgrau, Palle (ed.). *Demonstratives*. Oxford: Oxford University Press, 1990.