



Thought without Representation

Author(s): John Perry and Simon Blackburn

Source: *Proceedings of the Aristotelian Society, Supplementary Volumes*, Vol. 60 (1986), pp. 137-151+153-166

Published by: [Blackwell Publishing](#) on behalf of [The Aristotelian Society](#)

Stable URL: <http://www.jstor.org/stable/4106900>

Accessed: 21/08/2011 05:23

Your use of the JSTOR archive indicates your acceptance of the Terms & Conditions of Use, available at <http://www.jstor.org/page/info/about/policies/terms.jsp>

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact support@jstor.org.



Blackwell Publishing and The Aristotelian Society are collaborating with JSTOR to digitize, preserve and extend access to *Proceedings of the Aristotelian Society, Supplementary Volumes*.

<http://www.jstor.org>

THOUGHT WITHOUT REPRESENTATION

John Perry and Simon Blackburn

I—John Perry

I see a cup of coffee in front of me. I reach out, pick it up, and drink from it. I must then have learned how far the cup was *from me*, and in what direction, for it is the position of the cup relative to me, and not its absolute position, that determines how I need to move my arm. But how can this be? I am not in the field of vision: no component of my visual experience is a perception of me. How then can this experience provide me with information about how objects are related to *me*?

One might suppose that while no component of my perception is of me, some component of the knowledge to which it gives rise must be. Perhaps I am able to infer where the cup is from me, because I know how things look, when they are a certain distance and direction from me. Without a component standing *for me*, how could this knowledge guide my action, so that it is suited to the distance the cup is *from me*?

But some philosophers think that our most primitive knowledge about ourselves lacks any such component: basic self-knowledge is intrinsically selfless. Something like this was presumably behind Lichtenberg's remark, that Descartes should have said 'It thinks' rather than 'I think'. And according to Moore, Wittgenstein approved of Lichtenberg's remark:

. . . the point on which he seemed most anxious to insist was that what we call 'having toothache' is what he called 'a primary experience' . . . ; and he said that what characterizes 'primary experience' is that in its case, "'I' does not denote a possessor'. In order to make clear what he meant by this he compared 'I have a toothache' with 'I see a red patch'; and said of what he called 'visual sensations' generally . . . that 'the idea of a person doesn't enter into the description of it, just as a (physical) eye doesn't enter into the description of what is seen'; and he said that similarly 'the idea of a person' doesn't enter into the description of 'having toothache' . . . he said that

'Just as no (physical) eye is involved in seeing, so no Ego is involved in thinking or having toothache'; and he quoted, with apparent approval, Lichtenberg's saying, 'Instead of "I think" we ought to say "It thinks" . . .'¹

I am sympathetic with Wittgenstein's view as I interpret it. There is a kind of self-knowledge, the most basic kind, that requires no concept or idea of oneself. The purpose of the present paper, however, is not to argue directly for this view, but to try to see how it could be so, by seeing how it is possible to have information about something without having any 'representation' of that thing. I begin by studying something a bit more open to view, the possibility of talking about something, without designating it.

I

It is a rainy Saturday morning in Palo Alto. I have plans for tennis. But my younger son looks out the window and says, 'It is raining'. I go back to sleep.

What my son said was true, because it was raining in Palo Alto. There were all sorts of places where it wasn't raining: it doesn't just rain or not, it rains in some places while not raining in others. In order to assign a truth-value to my son's statement, as I just did, I needed a place. But no component of his statement stood for a place. The verb 'raining' supplied the relation *rains* (*t, p*)—a dyadic relation between times and places, as we have just noted. The tensed auxiliary 'is' supplies a time, the time at which the statement was made. 'It' doesn't supply anything, but is just syntactic filler.² So Palo Alto is a constituent of the content of my son's remark, which no component of his statement designated; it is an *unarticulated* constituent. Where did it come from?

In approaching this question, I shall make five initial assumptions, which together will provide a framework for analysis. First, I shall assume that the meaning of a declarative sentence *S* can be explained in terms of a relation between uses of

¹ G. E. Moore, *Philosophical Papers*, New York: Collier Books, 1962, pp. 302-3.

² Note that if we took 'It' to be something like an indexical that stood for the location of the speaker, we would expect 'It is raining here' to be redundant and 'It is raining in Cincinnati but not here' to be inconsistent.

S and what is said by those uses—the propositional content of the statement made. Consider the declarative sentence ‘I am sitting’. Different people at different times say quite different things by using this sentence. What they say depends in a systematic way on the *context*—the facts about the use. The pertinent facts in this case are the user and the time of use. An explanation of the meaning of ‘I am sitting’ quite naturally takes the form of a relational condition:

A use u of ‘I am sitting’ expresses a proposition P iff there is an individual a and a time t such that

- i) a is the speaker of u ;
- ii) t is the time of u ;
- iii) P is the proposition that a sits at t .

The second assumption is that the propositions expressed by statements—at least the simple sorts of statements we shall consider here—have *constituents*. Their constituents are the objects (relations, individuals, times, places, etc.) that they are about. Thus the constituents of my statement that I am sitting are me, the present moment, and the relation of sitting. The notion of a constituent, while intuitive, flies in the face of a long philosophical and technical tradition. I will not try to justify the notion here, but simply hope that skeptical readers can make enough sense of it to find what follows thought-provoking.

The third assumption is that a declarative sentence has significant components, the meanings of which can be explained in terms of the relations between uses of these components and the objects those uses stand for or designate. Let us suppose that in our sentence, the components are the three words, ‘I’, ‘am’, and ‘sitting’. We can explain their meanings as follows:

- A use u of ‘I’ designates an object a , iff a uses ‘I’ in u ;
- A use u of ‘am’ designates a time t , iff t is the time at which u occurs;
- A use u of ‘sitting’ designates a relation R , iff R is the relation *sits* (a , t).

In the first two cases, facts about the use affect the object designated. This is not so in the third case; no variable for the use appears on the right of the ‘iff’. Expressions of the first sort we call ‘context-sensitive’; those of the second we call ‘context-

insensitive', or 'eternal'. In this example, each of the components is a separate word, but this is not necessary, and isn't even plausible in the case of this simple sentence. A more plausible syntactic analysis would also find the component verb phrase 'is sitting'. This we could take to designate a more complex object, say a propositional function:³

A use u of 'is sitting' designates a propositional function $P(x)$, iff there are u' , u'' , R , and t such that

- i) u' is a use of 'is' that designates t , and u' is the initial part of u ;
- ii) u'' is a use of 'sitting' that designates R , and u'' is the second part of u ;
- iii) for any a , $P(a)$ is the proposition that $R(a, t)$.

The fourth assumption is that the meaning of a sentence is systematically related to the meanings of its components. In the simple example I have given, we can see what the relationship is (ignoring the verb phrase, for simplicity):

A use u of 'I am sitting' expresses the proposition P iff there are u' , u'' , u''' , a , t , and R such that:

- i) u' is a use of 'I' that designates a ;
- ii) u'' is a use of 'am' that designates t ;
- iii) u''' is a use of 'sitting' that designates R ;
- iv) u consists of u' followed by u'' followed by u''' ;
- v) P is the proposition that $R(a, P)$.

The fifth assumption is that a statement made by the use of a sentence is true, just in case the proposition the statement expresses is true.

The picture presented by this approach suggests a principle, which I shall call *homomorphic representation*:

Each constituent of the proposition expressed by a statement is designated by a component of the statement.

It is this principle to which my son's remark is counterexample. The propositional content of his use of 'It is raining' was that it was raining, at that time, in Palo Alto. But no component of his statement designated Palo Alto.

³ That is, a function whose values are propositions, not one whose arguments are, as the phrase might suggest to those outside philosophy.

II

We saw that there were basically two ways in which an articulated constituent is supplied. It can be built into the meaning of the expressions that it supplies a given constituent in any context of use, as we supposed to be the case with *sitting*. Or the meaning can simply identify a certain relationship to the speaker, a role that different objects might play, in different contexts of use. In the case of *I* the relationship is that of identity.

I suggest that unarticulated constituents are also supplied in these two ways. They can be fixed by meaning, once and for all, or the meaning may just fix a certain relationship, that the unarticulated constituent has to the speaker. This is, we can have eternal, and context-sensitive unarticulated constituents.

To this remark, one might reasonably ask what meaning it is, that either fixes the unarticulated constituent, or fixes the relationship it has to the speaker. After all, the problem is that there is no component of the sentence that designates the unarticulated constituent; hence, it seems inappropriate to begin by dividing the ways that it gets designated.

The unarticulated constituent is not designated by any part of the statement, but it is identified by the statement as a whole. The statement is *about* the unarticulated constituent, as well as the articulated ones. So, the theory is (i) some sentences are such that statements made with them are about unarticulated constituents, (ii) among those that are, the meaning of some requires statements made with them to be about a fixed constituent, no matter what the context, while (iii) others are about a constituent with a certain relationship to the speaker, the context of use determining which object has that relationship.

'It is raining' clearly has a meaning of the second sort. Let's assume, for a moment, that the unarticulated constituent for any use of this sentence is simply the place, at which the use occurs. Then an analysis of its meaning would be:

A use u of 'It is raining' expresses a proposition P , iff there are u' , u'' , u''' , t , P , and R such that

- i) u' is a use of 'It';
- ii) u'' is a use of 'is' that designates t ;
- iii) u''' is a use of 'raining' that designates R ;

- iv)* u occurs at P ;
- v)* u consists of u' followed by u'' followed by u''' ;
- vi)* P is the proposition that $R(p, t)$.

Clause *iv)* pertains to the unarticulated constituent. Unlike clauses *ii)* and *iii)*, it does not pick up a constituent designated by a component, but simply goes straight to the context, in this case, the facts about where u occurred.

It will be useful to have a term for that part of the context which determines the unarticulated constituent. I shall use the term 'background' for this. The background facts in this case are those about the location of the statements.

An analysis of 'It is raining here' would differ, just that instead of clause *iv)* we would have:

- iv)* u'''' is a use of 'here' that designates p

(with the rest of the condition changed as necessary to accommodate u''''). The place would then be an articulated rather than an unarticulated constituent of the proposition.

The supposition that 'It is raining' simply leaves unarticulated what 'It is raining here' articulates is not very plausible, however. Suppose, for example, that my son has just talked to my older son in Murdock on the telephone, and is responding to my question, 'How are things there?'. Then his remark would not be about Palo Alto, but about Murdock. All we should probably say as part of our analysis of the meaning of 'It is raining' is simply:

- iv)* u is about p .

This is not to deny, of course, that a good deal more could be said concerning the factors that determine which places a use of this sentence is about. The intentions and beliefs of the speaker are clearly key factors. My son's belief was about Murdock, and his intention was to induce a belief in me that was about Murdock by saying something about Murdock. Here it is natural to think that we are explaining which unarticulated constituent a statement is about, in terms of something like the *articulated* constituents of the beliefs and intentions it expresses.

My example of context-free provision of an unarticulated constituent is somewhat fanciful. Suppose there is a dialect,

spoken only by very chauvinistic San Franciscans. In this dialect, the sentence 'It is raining' is used to state the proposition that it is raining, at the moment of utterance, in San Francisco. ('It is raining here' is used for other locales the speakers of this dialect might find themselves in.) This is the proposition a speaker of this dialect asserts with 'It is raining' no matter where in the world it is spoken. San Francisco is then an unarticulated constituent of the propositions expressed by statements using this sentence. It is determined in a context-insensitive way.

III

Simpleminded as it is, this little theory establishes, I think, that there is no basic problem with a statement being about unarticulated constituents. In particular, we do not need first to find an expression, hidden in the 'deep structure' or somewhere else, and then do the semantics of the statement augmented by the hidden expression. Things are intelligible just as they appear on the surface, and the explanation we might ordinarily give in non-philosophical moments, that we simply understand what the statement is about, is essentially correct.

Still, it might seem that to correctly use and understand statements with unarticulated constituents, we must have, or be able to provide, expressions that designate them. When I hear my son say 'It is raining', and learn thereby that it is raining in Palo Alto, it seems I must have understood that his remark was about Palo Alto. And to do this, it seems I must have in my mind some concept or idea of Palo Alto, with which I can identify it as the right place. And as we noted, it seems that what made his remark about the weather in Palo Alto, in one case, and about the weather in Murdock, in the other, was his intentions and beliefs—what he had in mind, as we might say.

I shall argue that this is not quite right, although not quite wrong, either. We can imagine linguistic practices that do not require their participants to have any way of articulating some of the constituents of the propositions we would take to be the content of their statements. The basic idea is that the unarticulated constituents earn their role in the interpretation of statements by their place in the role of the thoughts that such statements express and give rise to, rather than by being designated by components of those thoughts. But once we have

imagined all of this, a slightly different way of handling things will suggest itself.

Consider a small isolated group, living in a place we call Z-land. Z-landers do not travel to or communicate with residents of other places, and they have no name for Z-land. When a Z-lander sees rain, he will say to others not in a position to look outdoors, 'It is raining'. His listeners then act appropriately to there being rain in Z-land: they close the windows in Z-land, cancel plans for Z-land picnics, and grab umbrellas before going into the Z-land out-of-doors. They have no other use for 'It is raining'. They do not call their sons in far-off places, or listen to the weather news, or read newspapers with national weather reports.

It would be natural to treat Z-landers' uses of the sentence *It is raining* as having Z-land as an unarticulated constituent. But what secures Z-land, rather than, say, San Francisco, as the unarticulated constituent of their discourse about rain? It is simply that the perceptions, that give rise to the beliefs that 'It is raining' expresses, are perceptions of the weather in Z-land, and the activities to which the belief gives rise are suited to rain in Z-land. Z-land is a constituent of the practice, or language game, in which the sentence 'It is raining' plays a role. There is no need to postulate a concept or idea of Z-land as a component of their thought in order to secure the connection to Z-land. The connection is secured by the role of the whole belief in their lives.

In the transaction we imagined with my son, there were three places that were relevant. First, there was the place his remark, my source of information, was about. Second, there was the place the belief I acquired from hearing him was about. Finally, there was that place rain in which would make appropriate the action to which my belief led me. As imagined, Palo Alto played all three roles. My son's remark was about the weather in Palo Alto, I took it this way, and going back to sleep was appropriate to rain in Palo Alto. But each of these connections might be broken. In a slightly different example, I would be misinterpreting a remark of my son's about rain in Murdock. His remark would be about one place, my belief about another. A little bit more elaborate change is required to break the second connection. Suppose we have spent the night in Sacramento, with the intention of driving back to Palo Alto early in the morning, so

we can play tennis. My son looks out the window, and says 'It is raining'. I take him, correctly, to be telling me about the weather where we are. But I have forgotten where we are. The action I take is appropriate to there being rain in Palo Alto, for if it were raining there, there would be no reason to leave early. But it is not appropriate to there being rain in Sacramento.

Given that we get information about the weather in various places, and have a repertoire of actions appropriate to weather in various places, our weather beliefs have a coordinating job to do, a job mine did satisfactorily in the original case, and unsatisfactorily in those we have just imagined. If our beliefs are successfully to guide our actions in light of the weather information we receive, they must reflect not only the kind of weather but also the place of the weather.

The Z-landers' beliefs have a simpler job to do. All of the information (or misinformation) they get about the weather, through observations or reports of others, is about Z-land. All of the actions they perform, in light of their weather-beliefs, take place in Z-land, and are appropriate or not, depending on the weather there. The connection between the place about which they receive weather information, and the place whose weather determines the appropriateness of their actions, is guaranteed by their life-style, and need not be coordinated by their beliefs.

Some psychologists and philosophers find it useful to postulate a 'language of thought', a system of internal representations, with a syntactic structure and a semantics, that is involved in belief, desire, and other mental activities and states. One goal of the present investigation is to develop concepts that will help us to understand the motives for attributing structure to thought, and the extent to which linguistic structure is the appropriate hypothesis. So I do not want to commit myself to any very determinate version of the language of thought. Still, we can use this hypothesis, bracketed, so to speak, to make the present point: there is no reason that thoughts that employ representations in the language of thought should not have unarticulated constituents, just as statements that employ sentences of natural language do.

IV

Still, it doesn't seem quite right to treat Z-landers' discourse

about weather just as we treated our own. A Z-lander semanticist would look at things differently. Having himself no concept of other places it might rain, he regards rain as a property of times, not a relation between times and places, as we do. He treats Z-landish discourse about the weather as homomorphic. What he provides as that which Z-landers believe and assert about the weather, the content of their discourse and thought, is something that to us seems to be but a function, from places to propositions.

There is something right about our Z-lander's point of view that we have not yet captured, and something right about ours that we do not want to lose sight of. There is some distortion in treating the Z-landers' uses of 'It is raining' just as we treat our own, as if there were a range of possibilities left open by their language that they simply fail to consider. Nevertheless, the possibilities we see, and they cannot yet express or think, are real.

Suppose we accept the Z-lander semanticist's opinion as to the objects of the Z-landers' attitudes—what they assert with a use of 'It is raining' and what they believe when they hear such a statement from a reliable source—but stick to our view of what those objects are. Then we would say that the Z-landers assert and believe propositional functions, rather than propositions. What would be wrong with this?

Let us back up for a moment. Beliefs have a semantic and a motivational or causal aspect: they are true or false, and they guide our action in achieving our goals. The two aspects are connected. The action to which a belief leads us, given our goals, should promote those goals if it is true. Thus my belief that it is raining in Palo Alto leads me to go back to bed, given my goal of sleeping late unless I can play tennis without getting wet. And if the belief is true, going back to bed will promote this goal.

Similarly, the Z-landers' beliefs about the weather lead them to actions that make sense if it is raining in Z-land. So, it seems that those beliefs ought to be true, depending on how the weather is in Z-land. And so it seems that the objects of the belief should be about Z-land, so that they will be true or false depending on the weather there. This last step leads us to attribute content to their beliefs non-homomorphically, for if we took the content to be a propositional function, rather than a proposition, it seems like

the connection between the semantic and the motivational aspects of their beliefs would be mysterious.

But this last step is not really necessary. There is another way to make Z-land relevant to truth of the Z-landers' assertions and beliefs. We can give up our fifth assumption, that a statement made by the use of a sentence is true, just in case the proposition the statement expresses is true. For the Z-landers' discourse about weather, a statement is true if the propositional function it expresses is true relative to Z-land. Z-land comes in not as an unarticulated constituent each Z-landish weather statement is about, but a global factor that all Z-land discourse about the weather concerns.

The point is to reflect, in our semantics, the lesser burden that is put on the Z-landers' assertions and beliefs compared to ours because of their impoverished sources of information and their limited repertoire of weather-sensitive actions. The only job of their assertions and beliefs concerning the weather is to deal with the nature of the weather in Z-land. Their assertions and beliefs are satisfactory, in so far as their 'weather constituent'—rain, snow, sleet, etc.—matches the weather in Z-land, whereas we need also to register the place of the weather. By taking the propositional content of their beliefs to be propositional functions, rather than complete propositions, and taking them to be true or false relative to Z-land, we mark this difference.

Let us develop a little more vocabulary to mark this distinction. We shall reserve 'about' for the relation between a statement and the constituents of its content, articulated and unarticulated. We shall say a belief or assertion *concerns* the objects that its truth is relative to. So the Z-landers' assertions and beliefs *concern* Z-land, but are not *about* Z-land.

V

As an alternative to this approach, we might consider taking Z-land to be a context-insensitive unarticulated constituent of Z-landish weather reports and beliefs. This would be plausible, in so far as it makes the relevance of Z-land a fact about the whole linguistic system, rather than about individual assertions and beliefs. It does not seem quite right, however. Suppose the Z-landers become nomads, slowly migrating westward. If their use of 'It is raining' is keyed to their new surroundings, we would

either have to say its meaning had changed, or that their reports were now false, whenever the weather in their new environs deviated from that in Z-land. Neither of these steps seems plausible. What we have contemplated is a change in their surroundings, not a change in the meanings of their sentences.

We can handle this under the approach of the last section, however. We can say that the place Z-landers' weather assertions and beliefs *concern* changes, as they move west. Or, if a schism develops, and different groups of Z-landers move off in different directions, severing connections with their old comrades, we can say that the different groups, though continuing to speak the same language, come to be concerned with different places. What is 'built into' Z-landish, at the current stage of its development, is that those who speak it are concerned with the weather where they are at, and their assertions and beliefs about the weather are true or false depending on the weather there.

VI

Could we apply this analysis to my younger son's remark? That is, could we interpret it homomorphically, taking it to express a propositional function, and say that it is true, because it concerns Palo Alto? But this would not be an accurate remark about English. Weather discourse in English does not uniformly concern the place the discussants are at.

Still, there is a little of the Z-lander in the most well-traveled of us. Talking on the phone and reading the national weather reports are one thing, talking to someone in the same room about the weather is a bit different. Our reaction to the local statement 'It is raining' is to grab an umbrella, or go back to bed. No articulation of the fact that the reporter's place and our place are the same is really necessary.

Something like the Z-landers' way of looking at things may be regarded as an aspect of our way of dealing with information about the weather, in circumstances in which the weather information we get is guaranteed either to be about or to concern our own location. And something like the semantics provided for the Z-landers' weather discourse is an aspect of the meaning of sentences like 'It is raining' in our language.

To borrow a phrase from Wittgenstein, we might say that the sentence 'It is raining' has a role in a number of different

language games. In those parts of our life where there is an external guarantee that the weather information we receive be about, and our actions concern, our own locale, there is no reason for our beliefs to play the internal coordinating role they need to at other times. When I look outside and see rain and grab an umbrella or go back to bed, a relatively true belief, concerning my present surroundings, will do as well as a more articulated one, about my present surroundings.

VII

There is a stronger point to be made, however. The weather in one's locale plays a special role in the life of humans. This is not necessarily the case for all agents that deal with information about the weather; the local weather of the National Weather Service Computer need have no special significance for it. But humans are affected in important ways by the weather around them, no matter where they happen to be. It is important that we be able to pick up information about the local weather perceptually, as we are able to do, and to act appropriately to it, by dressing warmly, taking an umbrella, or grabbing the sun-tan oil, as the case may be. These actions which help us deal with the local weather need to be under the control of beliefs that are formed through perception of the local weather. Efficiency suggests that there should be states of belief, typically caused by observations of the weather around one, and typically causing behaviour appropriate to that weather. That is, there should be a belief state⁴ that intervenes between perception of rain and behavior appropriate to that weather. That is, there should be are required to be *about* the place of the believer, then they must differ from person to person, depending on where they are, and even in a mobile individual, from time to time. Those in Phoenix should have their rain-behavior controlled by beliefs about Phoenix, those in Palo Alto should have their rain-behavior controlled by beliefs about Palo Alto, and so forth.

This could happen in two ways. One is that those belief states

⁴ The term 'belief state' suggests to many the total doxastic state of the agent, but I do not use it in that way. Two agents, each of whom has just looked outdoors and seen rain, could be in the same belief state, in my sense, in virtue of the common aspect of their total states that would lead each of them to say, 'It is raining', even though there is little else they would both be disposed to say.

that directly control behavior for local weather merely *concern* local weather, rather than being about it. All believers who had just seen rain and were about to open their umbrellas would be reckoned as believing the same propositional function, but the truth-conditions of their beliefs would differ with their location. The other would be to have these belief states correspond to a sentence like 'It is raining here'. This sentence makes a statement about the local weather, no matter who says it and where; an analogous belief state would be about the local weather, no matter who was in it and where. On this view, the believers would be in the same state, but would not believe the same thing, because the state contains an 'indexical' component.

We need both alternatives. An internal 'indexical' component of weather beliefs, that makes them *about* the weather in one's locale, is not necessary to understand beliefs with the causal role we have envisaged, intervening between local observations and actions appropriate to local conditions. It suffices that one's beliefs *concern* the local weather. Furthermore, using the indexical correctly is the same sort of ability as grabbing an umbrella when one sees rain. 'It is raining here' is an assertion appropriate when one sees rain, no matter where one is.

But a state corresponding to 'It is raining here' also has an important role to play for those who have access to information about weather in various places, and reason to communicate facts about their own local weather to others elsewhere that have such access. Such a state is best conceived as one which can be nomically tied to beliefs *concerning* the local weather and non-nomically tied, via beliefs about one's location, to beliefs *about* the local weather. I hear on the radio, 'It is raining in Palo Alto'. I believe that it is raining here, for I know that I am in Palo Alto. As a result I believe that it is raining, a belief at a more primitive level, that concerns Palo Alto. As a result, I get my umbrella.

The suggestion is, then, that our beliefs about the weather have a certain structure. At the bottom there are what we might call 'primary beliefs' about the weather, which are like the Z-landers' beliefs. These concern the local weather, and are true or false depending on it. They are typically caused by observations of local weather, and typically lead to action appropriate to local conditions. This is all our hypothetical Z-landers have, perhaps all that children have at certain stages of development,

and often all that we need. Above these are indexical beliefs, which are *about* the place that the more primitive beliefs merely *concern*: It is raining here.

At the top are beliefs that correspond to more sophisticated forms of getting information about the weather: reading or listening to news-reports, talking on the phone, and so forth. These beliefs are about various places, in virtue of relatively context-insensitive components of belief: It is raining in Palo Alto, It is raining in Murdock, and so forth. At the middle level are identificatory beliefs, that allow information at the top level to be translated into action at the bottom level: this place is Palo Alto.

VIII

This all suggests, I hope, a possible approach to the problem sketched at the beginning. What each of us gets from perception may be regarded as information concerning ourselves, to explain connections between perception and action. There is no need for a self-referring component of our belief, no need for an idea or representation of ourselves. When a ball comes at me, I duck; when a milkshake is put in front of me, I advance. The eyes that see and the torso or legs that move are parts of the same more or less integrated body. And this fact, external to the belief, supplies the needed coordination. The belief need only have the burden of registering differences in my environment, and not the burden of identifying the person about whose relation to the environment perception gives information with the person whose action it guides.

Lichtenberg's original remark was that one should say "There is thinking", just as one says "There is lightning".⁵ I have picked a somewhat less dramatic type of weather to serve as an analogy to self-knowledge, and developed it at somewhat greater length. Such analogies can carry us only so far, of course, but that is as far as I shall try to go in this paper.⁶

⁵ See Georg Henrik von Wright, 'Lichtenberg', in Paul Edwards (ed.) *The Encyclopedia of Philosophy*, New York: The Macmillan Company, Reprint Edition, 1972, volume 4, p. 464.

⁶ Recognition of the need for a distinction between what I here call *concerning* and *being about*, and the necessity to investigate non-homomorphic representation, was forced upon me by Joseph Almog and Bob Moore in the course of conversations about the motivation for propositions with truth values relative to times, as found in David Kaplan's work on demonstratives. The present approach is the result of conversations with Jon Barwise, David Israel, Bob Moore, John Etchemendy, and others.

THOUGHT WITHOUT REPRESENTATION

John Perry and Simon Blackburn

II—Simon Blackburn

WHAT ABOUT ME?

I

Perry's Strategy

John Perry introduces his suggestive paper by reminding us of a puzzle connected with reference to oneself: how is it possible that a simple experience should provide me with information about how objects are related to me, when no component of the experience is a perception of me? He uses this question to suggest sympathy with a possibly Wittgenstein view: 'there is a kind of self-knowledge, the most basic kind, that requires no concept or idea of oneself'. And at the end of the paper, the suggestion is that whilst 'what each of us gets from perception may be regarded as information concerning ourselves, to explain connections between perception and action', nevertheless 'there is no need for a self-referring component of our belief, no need for an idea or representation of ourselves'.¹

It is, I think, a little unclear how this sympathy with non-representative self-government (the Wittgensteinian or Lichtenbergian view) is motivated by the original puzzle. For that puzzle might be answered as, for example, Evans answers it. Evans similarly asks how we can have knowledge 'of a state of affairs which involves a substantial and persisting self, simply by being aware of (still worse, by merely appearing to be aware of) a state of the world'². Agreeing that we cannot get something for nothing, yet that nothing more than perception of (say) a tree is called for from the perceptual side, he answers the puzzle by requiring that the perceptual state, if it does sustain the thought 'I am seeing a tree', must occur in the context of certain kinds of knowledge and understanding. These are, in his view, essentially referential: I can think 'I am seeing a tree' because I can

¹ John Perry, 'Thought without Representation', p. 137 (all subsequent references to Perry are to this paper, in this volume).

² Gareth Evans, *The Varieties of Reference*, Oxford University Press, 1982, p. 231.

conceive of myself as of the kind which I envisage when I envisage someone seeing a tree—‘that is to say, a persisting subject of experience, located in space and time’. The persisting subject can equally be ascribed tensed properties—I was seeing a tree, or I will see a tree. This requires a conception of oneself as a located and enduring thing, and possession of this conception requires mental capacities which go beyond anything found in the single perception, or in the simplest cases of keying of perception to behaviour. But they provide the context of capacities and dispositions which create self-consciousness.

Evans’s solution directly contradicts the Wittgensteinian position, for there is full scale involvement of a concept or idea of oneself, yet it appears a possible and indeed attractive view. We shall see later if there is good reason for modifying it. It may be, for instance, that it takes too little account of the elusive nature of the self, the difficulty of making subject into an object of experience, and so on. Certainly, however, Evans must be right in remarking that ‘it is not a good idea, in attempting to determine the content of a person’s judgement, to examine nothing but the content of the perceptions which can legitimately give rise to it.’³

John Perry appeals to the general semantic framework of his paper to introduce the possibility of talking about something without designating it. In turn he uses that to introduce the possibility of having information about something, without having any representation of that thing. I am sufficiently sympathetic to the framework, with its straightforward use of propositions and their constituents (which they are about), of sentences and their significant components, not to raise any general doubts at the outset. But ‘proposition’ is a term of art, and we need to be careful about the constraints on its use. The main connection which requires care is that between the proposition asserted, and whatever is *thought* or *understood* by the utterer or his audience.

Thus Perry asserts that his son’s remark ‘it is raining’ expresses a proposition one of whose constituents is the place at which it was made—Palo Alto. Now it is certainly right that such a remark, made at that place, is made true or false by the

³ Evans, p. 233.

weather at Palo Alto, and may quite properly be said to be *about* the weather at Palo Alto. Since nothing in the sentence designates Palo Alto, the place becomes, in Perry's terms, an unarticulated constituent of the proposition. Now if we say this, we naturally have to sever some other attractive connections—particularly between identifying the proposition made by a remark, and understanding it. For the metaphor of a 'constituent' of a proposition must permit this principle:

You can identify a proposition only if you know which each of its constituents is.

It also seems reasonable to suppose:

You can identify your whereabouts (or that of a speaker) as Palo Alto only if you know that that is where you are (or he is).

But now we face the fact that Perry's son's remark can be understood, and indeed verified, by a hearer who has no idea that he is at Palo Alto, or who would deny that he is there, or assert that it is sunny at Palo Alto. Any hearer who is lost or under a misapprehension as to where he is could still fully understand the remark. People lost know what they express by saying 'it's foggy here'—that might be why they are lost, perhaps. 'It's raining here' is just the kind of thing of which the disorientated prisoner with no idea of his whereabouts might be left aware. This is no objection to Perry's notion of 'aboutness', nor to his conception of the proposition. However, what we are then forced to say on this conception, in the light of my two principles, is that you can understand what someone said without knowing which proposition they expressed. There is nothing impossible about that—but the notion of proposition so introduced may be of less use in connexion with Perry's original enterprise than might have been expected. That enterprise, remember, is to approach the topic of thought without representation. Now the notion of a proposition which is at this much distance from understanding may be a good notion with which to think about *information* and *truth*. But it is evidently not quite the notion—or at least not evidently quite the notion—with which to think about understanding. And thought, surely, goes with understanding.

We might introduce a distinction at this point, between what a speaker's remark is about in Perry's sense, and what the speaker must knowingly denote, this being what he must know his remark to be about if he makes the remark in full understanding, or equally what an audience must know the remark to be about if it receives it with full understanding. A speaker saying 'it's raining here' does not knowingly denote Palo Alto, even when he is there, for as we have seen, it is not a requirement on full understanding of that sentence uttered at that place that he should know where he is. What he must know is that rain infuses the portion of public space that he currently occupies—whichever that may be. And this is what a competent audience must take from the remark. This is why the audience understands a speaker saying 'it's raining here' when hearing him over a telephone, and not knowing where he is ('here' differs from 'this', used to effect demonstrative reference, in this respect, for with that it does seem required that the hearer know what was the object of the demonstrative). 'Knowingly denote' is of course highly intensional, for a speaker knowingly denoting the region of space in which he is currently located need not be knowingly denoting Palo Alto, even if that is where he is (one can think of cases in which if the speaker knew he were at Palo Alto, he would revise the opinion that it is raining—he might think it fallout from garden sprinklers, there, for instance).

It is not too surprising if a speaker, or hearer, need have no representation of the thing (place, time) that a remark is about, in Perry's sense, since to understand the remark they need no awareness or knowledge that it concerns that object, or place or time. Perhaps this is most obvious if we take the example of a people lacking the ability to locate times in an objective time order. 'It just climbed in there' is said and heard with understanding only by those who know that the time of the climbing is said to be just before that of the saying, but even if the time of the saying was midnight Greenwich Mean Time, this need not figure in their thoughts, nor be capable of doing so. They might have the capacity to think 'it just climbed in there' without any capacity to express or judge a proposition of the form 'it climbed in there at time t '.

Perry's aim would be better furthered if it could be argued that *knowing* denotation requires no representation, so that a

speaker might knowingly denote the region of space he is in without having any kind of representation of it. But it is hard to see what would support that. Certainly the ordinary information channels are naturally thought of as representing my spatial surroundings to me, and indeed the directness with which they do that—the unreflecting awareness that things are over there, near my hand, to the left or right—probably provides our best paradigm of ‘representation’. Now a speaker might understandingly say ‘I will wait for something to happen here’ (for instance) when he is in one of those tepid, anaesthetized, information-free, states in which no information from his surroundings is being received. It might be suggested that this is a case of thought without representation—since there is no presentation of the surroundings, there is no representation of them either, yet the thought concerns the surroundings, and the speaker knowingly denotes them. But this is unconvincing, for the subject retains his capacity to accept information from his surroundings—a squeak just behind or a glimmer a few feet to the right—and can imagine to himself the hoped-for display. His capacity to figure himself as being at one point in space is unimpaired. So the better thing to say is that he can represent to himself his surroundings, although he lacks information about them. Again, then, we have no clear motivation for allowing thought without representation.

Is the situation changed when we consider thought of a relatively impoverished kind, as Perry introduces with his Z-landers? It is important to the case that Z-landers are not just indifferent to the weather elsewhere, so that they find no need to signal that it is their own place they are talking about (it goes without saying). If this were the situation, then we would have an instance of a quite general kind of ellipsis, in which an intended value for a variable need not be spelled out because everyone will take it that one particular thing or kind of thing forms the intended value. This happens when we describe something as dangerous—suppressing explicit mention of the value of the implicit variable (to us, or to me). Z-landers by contrast are people who cannot ‘express or think’⁴ the possibility of rain elsewhere than where they are. The crucial idea is that

⁴ Perry, p. 146.

'the connection between the place about which they receive weather information, and the place whose weather determines the appropriateness of their actions, is guaranteed by their life style, and need not be coordinated by their beliefs'.⁵ Certainly, to live out their lives as Perry describes, these people need no concept of Z-land itself being one place among others, nor of rain elsewhere.

II

The semantics

Perry wants his semantics to reflect the 'lesser burden' put on the Z-landers' assertions and beliefs, compared with ours. By saying that the object of their belief is a propositional function true relative to Z-land (so that Z-land is involved only externally to the objects of their belief) he hopes to achieve this, but at the cost of failing to allow them self-contained, truth evaluable, beliefs at all. This is a clear cost, but it may be queried whether the proposal does, in fact, achieve its end in any case. For the Z-landers' incapacity to understand that their place is one amongst others, in any of which it may be raining, seems poorly captured by actually attributing to them cognition of a function 'it raining at x ' evaluated relative to Z-land. The proposal puts Z-land outside the sphere of their cognition, but it leaves in there something which should not be—namely understanding of a general property (it raining at a place), which introduces exactly the possibilities which they cannot 'express or think'.

Consider by analogy the following case. People can be kind to people, say, and also kind to animals. Now imagine a culture to whom the very idea of kindness to animals is foreign. It never occurs to them to think in those terms—perhaps they have a philosophical or religious tradition which denies animals sentience, for example. We do have a problem when we try to represent their sayings. It would seem wrong to translate 'Genghis is kind' homomorphically, without noticing the divergence of domain; it might seem wrong to translate it as 'Genghis is kind to people' if, for instance, the limitation to people is one of which the culture betrays no awareness (no explicit restriction of domain enters into their thoughts, so we

⁵ Perry, p. 145.

should not translate them as if it does). But it would seem equally wrong, for the very same reason, to translate it as 'Genghis is kind to x ' conceived of only as a propositional function, true relative to people. If anything, that suggests the very idea to be avoided, that they are involved with a generally applicable concept, with a contingent limitation of its interesting domain. The principle transgressed seems universally desirable: if the possibilities they understand are limited in some definite way then it is wrong to attribute concepts to them whose application is not thus limited.

Perry notices, of course, the discomfort inherent in the idea that what is on the surface a self-standing proposition should come out as only a propositional function. He dismisses the alternative of calling Z-land a context insensitive, unarticulated component of the reports and beliefs because he thinks we then need to say either that the meaning of their remarks changes, or that their reports become false, if a slow, unnoticed migration takes place.⁶ These alternatives can, however, be avoided even within the spirit of this semantics. Context is introduced simply as encompassing any facts about the use of a sentence.⁷ So there is a sense in which Z-landers' remarks are context insensitive—there is no fact about day-to-day use which introduces *variation* of a place component in different remarks. Call this 'context variation'. There is no day-to-day, or individual context variation in their weather remarks in the way that there is with ours. But there is another more important sense in which their remarks are nevertheless context sensitive. For given the way they work, it simply *is* a fact about the use which makes it so that Z-land is either the object of their remarks or in any other way concerned in them. Perry brings Z-land into the teleology and the normative aspect of their states: he points out that the Z-landers' beliefs about the weather lead to actions that make sense if it is raining in Z-land, or that their beliefs ought to be true 'depending on how the weather is in Z-land'.⁸ But these dependencies are derivative. The fact about the use on which they depend is that Z-land is where these people are. Given this we could say that Z-land is a context dependent, unarticulated

⁶ Perry, p. 147–8.

⁷ Perry, p. 139.

⁸ Perry, p. 146.

component of all their weather reports and beliefs, and incur no burden of saying that their reports change meaning or become false as they migrate, nor of seeing the remarks as context variable, as ours are. This would maintain what is right in Perry's theory without supposing that Z-landers express anything less than full beliefs in their utterances.

There is however a much more important aspect of this. Perry rightly contrasts the states which need to be postulated to explain appropriate perception linked behaviour, with those involved in full use of an indexical. Whatever else it needs, an animal needs its behaviour to be appropriately controlled by its perceptions, and for that link to be working well it needs no conception of itself as occupying one particular place amongst others. An animal could be good at behaving appropriately—taking shelter—when it perceives rain, without having any conception of it as raining here. But is it right to see this as related to the contrast between having beliefs that *concern* a place, as opposed to having beliefs *about* that place? Perry seems to connect the sophistication of our context variable use of indexicals (justifying the 'about' interpretation) with our need to interpret communications coming from different places ('listening to news reports, talking on the phone, and so forth'). 'Here' thus comes in as a term whose role is to coordinate communication which, as it happens, can involve speaker and hearer being in different places. 'Concerning', by contrast, with no full scale use of the indexical, is appropriate where this kind of coordination problem is somehow out of court.⁹

But the fundamental distinction is surely not between context variation and context dependence, nor between 'about' and 'concerning'. It is not to be drawn in terms which, as we have seen, make no immediate reference to the understanding or conceptual repertoire of the subject. The real question is whether the subject possesses the concept of an objective spatially arrayed world, in which its own standpoint is that from one particular place. If it has such a conception, it can think the difference between, for instance, it raining here, and it raining where I am (it might rain here without raining where I am), or between it having rained here recently, and it having rained on

⁹ Perry, p. 151.

me recently, or it being about to rain here, and it being about to rain on me. The indexical for place is used in assertions whose understanding requires that comprehension. It is its comprehension of this range of possibilities, rather than the need to coordinate communications from distant places, which is fundamental. (If the subject did not possess the scheme to begin with, it would face no problem of coordinating communication about different places.)

Perry is right that the most fundamental level of thought, in which when a ball comes at me I duck or when a milkshake is put in front of me I advance (or retreat), requires no idea or representation of ourselves. It is facts external to the belief—facts about the integration of our control systems—which, as he puts it, supply the needed coordination. There need be no self awareness, and no self knowledge, because there is here no exercise of the capacities which define self-consciousness. At least, this is so if these cases approximate to simple reflexes, such as the movement of the eye to fix directly sudden movement registering on non-central parts of the retina (ducking is more like this than lunging for a milkshake).

Perry does however suggest that even at the fundamental level, what we get from perception may be regarded as information *concerning* ourselves, to explain the connections between perception and action.¹⁰ I am not clear how this works. If the integration of my control systems does the explaining, it is not evident why we need a reference to myself in the identification of any belief state, even if the reference is external, to be couched in terms of 'concerning' rather than 'about'. In the third person: he believes, concerning himself, *x* is in front of a milkshake, seems to play a redundant, or, in its suggestion of there being an identity judgement in the offing, a potentially misleading role, compared with: he is aware of a milkshake. The misleading implication is that there is something common to the ordinary case and a case in which, having thought that 'being in front of a milkshake' has an instance the subject goes on to think or act in ways appropriate to thinking 'Lo!, it is I!' One can set up such two-component cases, but only in complete contrast to normal perception.

¹⁰ Perry, p. 151.

In this section I have not had much good to say about the distinction between 'concerning' and 'about'. But I do not intend to dismiss it as a proposed research strategy. It would certainly tie in with the attempt to define a non-representative conception of the self if we could go on to construe 'belief, concerning oneself, that x is F ' as itself involving no pronoun whose function is ultimately explained only by involving *bona fide* reference. If, for instance, instead of 'concerning oneself' one could substitute some special, adverbial notion ('ego-centric' belief that x is F) we might move in that direction. But the pros and cons of that course have not, I think, yet been established.

III

Thought without Representation?

For Lichtenberg or Wittgenstein to appeal, it must remain true, when we turn from the level of simple self-government to that involved in the ordinary comprehension of the indexicals, that we require no representation or idea of the self. But is this right? In understanding myself to be here (when I might have been there) I know myself to be a kind of thing with a location in the world, which moves, acts, perceives; I have a history, I am quite like you, I will die. I can represent myself to myself, as a human being or, in short, a living animal, with a shape and an age.

True, I can think those other things which seem to throw open the kind of thing I am and which tempt us to a Cartesian conception of the real self. I can think that I might have lived earlier or later, or that it is an exercise of luck or a miracle that I have this perspective on the world and not that of someone else. I can wonder whether reality might be entirely my mental construct. I can wonder what it would have been like for me if my parents had never met, and if I am told that it would have been like nothing for me, I can still think of that as at best a kind of accident, as if I might have popped up somewhere else in any event. Faced with the personal identity puzzle cases, I can insist on trying to determine what it would be like for *me* after the sinister events, even when I know that all objective facts leave no single reasonable answer (and wondering whether it will be I

who wakes up in the red room is not wondering whether something as good as its being me is going to happen).¹¹ If we trusted these thought experiments, we would indeed lose any representation of ourselves, for we shrink to pure Cartesian subjects, extensionless and without essence. But is there any reason to trust them? At least a first reaction must be to see if we can diagnose them, by explaining why our imaginative powers give rise to what are only illusions of possibility.

In the passage Perry quotes, the Wittgensteinian suggestion appears to be that the word 'I' loses its referential role when we talk of mental attributions, rather than of physical ones ('I intend to go walking', as opposed to 'I weigh eleven stone'). But a better way to avoid the illusions can carve the uses differently. Thus it seems that there is something right about the thought that in certain contexts 'I' can be used—even if in the light of more thought educated usage might abandon it—although it does not function as a referring expression. Bernard Williams locates the suspicious cases. There is, he points out, a way—in fact, two ways—of imagining 'myself being a racing driver' which justifies the inclusion of reference to me: I 'am prepared, as it were, to accept a lot of my actual self in the fantasied scene'.¹² The two ways involve firstly participation (I grip my typewriter, perhaps go 'brrrrm, brrrrm', envisage cinders flying etc.) and in the second way an external view of me—the real me—crowned with garlands, sprayed with champagne, a hero to my children etc. But there is another mental process equally easily described with 'me' in it, which does not sustain the idea that reference to myself is really involved. This is the process which leads to thoughts of transference ('I might have been Napoleon') and which led Schlick to say that he could imagine himself seeing his own funeral. And as Williams says, all that is really reported in this way is Schlick's visualisation of his own funeral. Similarly I might report myself as imagining 'myself being Napoleon' when what I do is imagine the desolation at

¹¹ Derek Parfit implies that the lesser question (whether it is as good as survival if. . .) is all that is properly left us if we abandon a Pure Ego theory which alone makes sense of these puzzles. I think, in what I believe to be a Kantian vein, that no such theory (no *a priori* psychology) is implicit in the way we make sense of the imaginings.

¹² Bernard Williams, 'Imagination and the Self' in *Problems of the Self*, Cambridge University Press, 1973, p. 39.

Austerlitz, as if I were present, viewing it ('vaguely aware of my short stature and my cocked hat . . .').¹³

According to Williams, and plausibly, such fantasies represent an enactment of the role of Napoleon: a mode of imagining which introduces no further rootless Cartesian 'me', but involves only the real me (fantasizing) and Napoleon (dead). If untutored language describes the fantasy as that of me being Napoleon, then the word 'me' is not referential in such a context—there is literally nothing transferred from the actual world to the imagined scene, in the way that Baker Street is transferred from the real world to Conan Doyle's fiction. Such imaginings reflect my power to adopt a different standpoint on the world in my thought. But this is just envisaging the world as it appears to someone in such-and-such a position. There is no 'I' who is transferred to that position.

It would be nice to think that the temptation to the Cartesian Self, or equally the alternative of supposing that 'I' has a non-referential but 'transcendental' and therefore mysterious employment, can be subdued by insisting on a hygienic redescription of the real content of the fantasy in these transference cases. But much more work would be needed to sustain this solution. One difficulty is this. The power of exercising the imagination in the 'transference' way seems to be a power close to that which reflective thinking of any complexity necessarily involves. An agent keyed to anything more than immediate stimuli needs to think out what the world would look like if he moved *there*, or had been there *then*. The defence of Wittgenstein and Lichtenberg will now perceive a crack. For if the imaginative possibilities can be best expressed without involving a referring 'I', then is there not the threat of that being true more generally? I originally introduced imaginings, as Williams does, as both especially prone to give rise to the illusion of the Pure Ego, and as forming a relatively isolated, manageable, area, in which the non-referential function of terms apparently referring to myself (or alternatively, the desirability of avoiding introducing myself into a true representation of the content of the imaginings) could be sustained. By contrast, in everyday contexts—in particular in the use of tensed first-person assertions—reference to a self meets

¹³ Williams, p. 43.

no obstacle, but in turn supports no Cartesian conception of the self referred to. This synthesis is thrown into doubt if the exercise of the imaginings is in effect an exercise of the very powers which allow me to tense predicates of myself, seeing myself in earlier or later events.

It must be replied that something more is to be found in real tensed self-predication—in particular the connection between genuine (first-person) thought and action. Envisaging that from the brow of the hill I will be seen by the hunter, I do not go there; imagining ‘myself’ as (someone who will be) at the top of the hill and seen by the hunter does not matter. But this suggestion leaves problems, for two reasons. Fantasies are only entertained, whereas predictions are believed, and this alone seems sufficient to explain a divergent impact on action, without *also* finding a difference in the content on the imagining. If there is a changed content, it must reveal itself when I merely entertain the genuine thought concerning myself, that I go to the top of the hill and get seen by the hunter, without giving it any particular credence, or letting it affect my actions. The other problem is that the connection between tensed predication and action does not obviously need to chime in with the idea of a shared object of reference. Even if ‘Of me it will be that . . .’ has a particular effect on action, it needs showing why that makes legitimate the transfer to ‘it will be that I . . .’.

In the course of defining his use of the ‘Generality Constraint’, whereby the ordinary referring status of ‘I’ is tied in with the capacity to make predications of myself as a thing present in other times and places, Evans remarks as an instance my grasp of the thought that ‘I was breast-fed, or that I was unhappy on my first birthday, or that I tossed and turned in my sleep last night . . . or that I shall die’.¹⁴ We do indeed grasp these thoughts. But does whatever is involved in that grasp sustain the idea of a common object of reference? If the thoughts equate

¹⁴ Evans, p. 209. The idea of course derives from Strawson’s *Individuals*, as Evans acknowledges (p. 103). As far as I can see Evans does no more than assert that ‘I’ thoughts do conform to the Generality Constraint in any kind of way that *shows* the ‘I’ to be referential. The rival can say that even if one can think ‘egocentrically’ of having been *F*, being *F*, being possibly *F*, going to be *F*, etc. this does nothing to suggest that egocentricity is a matter of a distinct reference. For even if permitting distinct predicates is a requirement on interpreting a term as referring, it does not follow that it is sufficient for it.

with 'this living animal' (I might thump myself here) 'was breast fed, or unhappy . . . and will die' then there is no problem about the common object of reference—the ordinary empirical self with arms dangling and feet together. But it is not so clear that we allow them to equate with that: back come the Cartesian thought experiments.

However, at this point the Wittgenstein-Lichtenberg comeback runs out of steam. For suppose the only good argument for saying that we should not allow the content of the tensed self-ascription to identify with the content of a sentence referring to the animal cites some variant of the transference thought experiments—'this animal might have had an unhappy birthday without my having done so'. Then in the present context they cut no ice, for we are entitled not to allow that this untutored expression of the content of these imaginings gives us any grip on the metaphysics of the self. The issue, remember, is whether the misleading, potentially non-referential uses of 'I' can be confined to those reports of imaginings in whose content I do not really figure. The suggestion is that by contrast there is a quite normal reference, to the living animal, in straight prediction and retrodiction about myself. If that is the agenda, it cannot be argued that such ordinary reference is never in place—because of the very transference 'possibilities' which are in the process of being quarantined.

Of course, other arguments may be waiting in the wings. But the standard arguments concern guaranteed reference, 'immunity to error through misidentification', and the possibility of self-reference without awareness of the features which will be essential to whichever animal I am. I agree with Evans that none of these works, which throws us back onto unreliable imagination.