

Location, Location, Location

Huw Price

Abstract: This piece was written as my Presidential Address at the Annual Conference of the Australasian Association of Philosophy, held at Melbourne University in July 1999. I discuss the view 'that we can't describe or theorise about the world from outside language.' I call this idea 'linguistic imprisonment', and take it to be a platitude, although one that is interpreted very differently by different philosophers. In so far as language does depend on contingencies of our own 'location', how should we theorise about such matters? I distinguish two approaches, called 'backgrounding' and 'foregrounding'. Roughly, the latter seeks to incorporate the contingencies into the content of claims that depend on them, whereas the former treats them as use conditions. I argue that linguistic imprisonment implies that not everything can be foregrounded, and apply the framework to a then-recent objection to expressivism by Jackson and Pettit.

Location, Location, Location

Huw Price

1. Introduction

It is a platitude in philosophy these days that we can't describe or theorise about the world from outside language. Let's call this the doctrine of linguistic imprisonment (being ready to plead irony, if necessary, if the confinement turns out to be less of a restriction than the penal metaphor suggests—if the prison turns out to be “false”).

Partly because it is a platitude, this doctrine has a status in contemporary philosophy analogous to that of the existence of medium sized dry goods. All sides lay claim to commonsense, and argue that their opponents fail to do it justice. For example, metaphysical realists claim that idealists mistake the commonsense doctrine of linguistic imprisonment for the absurd view that language somehow “constructs” or gives rise to the world. For their part, idealists object that it is these same realists who fail to respect the doctrine, by taking seriously issues which in fact it disallows—issues about the “real” reference of terms, for example, where this is taken in a sense which could only make sense from “outside” language. Roughly speaking, then, the charges are that idealists overestimate and metaphysical realists underestimate the significance of our “imprisonment” in language.

My sympathies are more with the idealists in this exchange. In this paper, however, my concerns are both more basic and somewhat tangential. I want to clarify the nature of our “confinement” in language, but I also want to defend the possibility of a kind of “imprisonment” which even my idealist allies are inclined to overlook: roughly, the possibility that the prison is multicellular—it contains many distinct cells, to some subset of which we may be confined.

But how much are we constrained by the prison of language? In one sense, surely, not at all—or no more than by our usual ability to do the impossible. (This is at least one sense in which language may be said to be a “false” prison.) However, this reassuring thought seems a little too swift. One disturbing thing about linguistic confinement is surely that it threatens us with unavoidable anthropocentricity, or relativity, or perspectivalism—with the inability to escape from a viewpoint which is in some sense “localised” to ourselves. This threat doesn't seem significantly less disturbing—to those disturbed by it at all—in the light of the conclusion that it is a matter of necessity.

Location, Location, Location

Huw Price

Abstract: This piece was written as my Presidential Address at the Annual Conference of the Australasian Association of Philosophy, held at Melbourne University in July 1999. I discuss the view 'that we can't describe or theorise about the world from outside language.' I call this idea 'linguistic imprisonment', and take it to be a platitude, although one that is interpreted very differently by different philosophers. In so far as language does depend on contingencies of our own 'location', how should we theorise about such matters? I distinguish two approaches, called 'backgrounding' and 'foregrounding'. Roughly, the latter seeks to incorporate the contingencies into the content of claims that depend on them, whereas the former treats them as use conditions. I argue that linguistic imprisonment implies that not everything can be foregrounded, and apply the framework to a then-recent objection to expressivism by Jackson and Pettit.

It might be argued, as in effect by Davidson, that the same considerations which show that the prison is false also show that the idea that it is local or relative is incoherent—that the very idea of alternative locations is a mistake. However, I don't think this quite meets the concern. The disturbing contingency need not be couched in terms of alternative viewpoints, at least in the sense which Davidson takes to be problematic. Thus, suppose we were to discover that our possession of certain groups of concepts depends on contingent features of our biological circumstances. In other words, we find ourselves saying: "If we had been different in these ways, we would not have had that group of concepts; we would not have been able to say the things we say with those concepts." This may still seem to make language unacceptably "local", even if for Davidsonian reasons we don't want to acknowledge that the creatures we would then have been would be language users at all, in the interpretative sense (according to which what counts as language is what can be interpreted from where we stand).

The concern here might seem relatively trivial. We are finite creatures, of limited cognitive capacities. As such, presumably, we can't talk about everything. Some things are off-limits to us on complexity grounds alone, perhaps. Why should it surprise or disturb us that had we been differently constructed—simpler or more complex, say—we would not then have talked about some of the things we do talk about, or would have talked about some things we don't in fact talk about? In some cases, surely, we would have talked about the same things but in different ways—the same objects, under different modes of presentation. Again, what should be surprising or disturbing about that?

In my view, these attempts to render imprisonment nonthreatening tend themselves to rely on a violation of the doctrine. In effect, they tend illicitly to take for granted a viewpoint external to our systems of concepts, from which it can meaningfully be said that our concepts refer to this or that subset of the domain of things in the world, or that this concept and that refer to the same objects in this domain. Later in the paper, I want to ask what remains of this reference-based "externalising" attempt to render imprisonment non-threatening, if we avoid an illicit violation of the doctrine. Initially, however, I want to make a case for taking linguistic contingency seriously—for regarding it as the basis of a non-trivial kind of linguistic imprisonment. For the present, then, I'll regard these referential externalist ways of down-playing the significance of linguistic contingency as off-limits. (I suspect that they would be off-limits for Davidsonians in any case—which means, I think, that Davidson's rejection of the notion of alternative conceptual schemes

cannot wholly lay to rest the disturbing consequences of the combination of imprisonment and contingency.)

Initially, I want to focus on the relationship between two things: on the one side, certain bits of language—concepts, or terms; on the other side, contingent features of language users on which the use of those bits of language depend. (I'll call these features of language users the contingent grounds of the use in question.) I want to do three main things:

- (i) to distinguish two different ways of thinking about this relationship between concepts and their contingent grounds;
- (ii) to point out that one way of thinking about it makes imprisonment more serious than the other way does; and
- (iii) to argue that the more serious kind of relationship is unavoidable—it is an empirical question how much it affects us, but we can't avoid it altogether.

2. Two distinctions

I need a couple of preliminary distinctions. These turn on two issues we can raise about our own linguistic practice.

Is linguistic competency hardware-specific?

In principle, perhaps, we might be “universal” linguists, capable of “running” or using any language at all. If so, then in so far as the language we actually run is special or distinctive, its distinctiveness is not necessitated by our physical circumstances. (It may be appropriate in virtue of our physical circumstances to talk about one thing rather than another, but it would be physically possible for us to do otherwise.) Alternatively, we might be running a language which is “hardware-specific” in various ways. Of course, these distinctions may cut differently at different linguistic levels—our language is obviously hardware-specific in certain phonetic respects, but might perhaps be hardware-independent in conceptual respects. I am interested in the conceptual level. Are our concepts hardware-dependent or hardware-independent? And if they are hardware-dependent in us, to what extent is this a feature of the concepts themselves, and to what extent does it merely reflect limitations in us, so that a “universal” linguist could run them without special hardware? (For present purposes I want to ignore hardware restrictions which are simply a matter of processing power. I'm also happy to be vague about what counts as hardware. The crucial thing is that it is not the sort of thing that we are free to change, by means available to normal speakers.)

Are concepts functionally homogeneous?

Is there significant functional “modularity” or “inhomogeneity” among our concepts? Is there any (non-obvious) sense in which they do different jobs, or serve different functions in our lives? Obviously we use different concepts to talk about different kinds of things, and it isn’t controversial that we may use such concepts for different purposes, in virtue of differences between the things talked about. As I’ve argued elsewhere, the less obvious and more interesting possibility is there might be cases in which the proper order of explanation goes the other way—cases in which we need to appeal to differences in the function of the talk, in order to explain apparent differences between the things talked about. (For example: the possibility that we might need to appeal to the distinctive function of normative concepts in our lives, in order to explain the differences, as they appear to us, between normative facts and non-normative facts. The motto of this approach is that function might explain phenomenology.)

Why does this question matter? Because it provides a potential non-hardware-specific sense in which we might turn out to be “isolated” or “located” by contingencies related to our use of language. Think of different conceptual modules or language games as like distinct programs. Now suppose that although we were “universal” linguistic machines, capable of running any possible program (given enough time and memory), we couldn’t run all possible programs at the same time. This would provide a concrete sense in which our conceptual standpoint was inevitably “located”, in some contingent way. (Here’s an analogy. Imagine a race of sporting all-rounders, capable of excelling in any game at all—Australians in the 1950s, perhaps. Obviously, not even such superbeings could participate in all sports at once. With one or two exceptions, the rules of the various games are simply incompatible with multitasking.)

Of course, another possibility is that functional modularity might be associated with hardware dependence. In that sense, it would not be a matter of choice which programs we ran. (Think of dolphins, who might be excellent at water polo, but useless at cricket.) But in principle language might be modular at the conceptual level without any corresponding modularity at a hardware level—a universal linguistic engine might simply go in for “multitasking”, running several programs at once, with different functions.¹

¹ However, this distinction might become blurred if we softened the notion of hardware, on the one hand, and recognised on the other that functions found in use in language are likely to have extralinguistic significance—in interesting cases there is likely to be a biological reason why we are running a particular program, and this might well amount to a hardware-dependency, broadly construed, even if the same wetware has the capacity to do other things.

Thus the two distinctions are independent, giving us a four-fold table of possibilities:

	Hardware-independent	Hardware-dependent
Functionally homogeneous	1	2
Functionally inhomogeneous	3	4

I am interested in the possibility that our own linguistic practice belongs in the fourth cell in this table, and in the philosophical consequences of this possibility. That cell seems to represent a more serious kind of imprisonment than the other three options. On the one hand, if language is functionally homogeneous then this excludes at least one kind of “locality” or “confinement”—that which would flow from the fact that actually we run some possible modules but not others. While on the other hand, if language is “all in software”, then any modularity is “non-vicious”—in principle, any speaker could run any module, even if not all at once. But the combination of multifunctionalism and hardware-dependence seems to leave us especially trapped. That’s what makes the fourth option particularly interesting.

In my view, it is ultimately an empirical issue whether our own constitution and the nature of language puts us in Cell Four. However, in order to show both that this is a genuine empirical possibility, and that it has interesting philosophical consequences, it is necessary to block two philosophical counterattacks. One of these counterattacks claims to establish a priori that we are not in Cell Four. It argues that it is always possible to interpret the content of concepts in a way which renders them non-hardware-dependent and unifunctional. The other counterattack claims to mitigate the philosophical consequences which would follow if we were in Cell Four. It argues that even if our concepts were hardware-dependent and modular, this would have no important consequences for metaphysics: it is the view mentioned earlier, that what matters is not the concepts or modes of presentation, but the world to which they refer, and that plurality and hardware-dependence in the former tells us nothing interesting about the latter.

This second counterattack is the “externalist” strategy I mentioned at the beginning. As I said there, I think that its appeal rests in part on a failure to take seriously the imprisonment constraint. Later in the paper, I want to ask what remains of this strategy if that mistake is avoided.

For the moment, then, I want to focus on the first counterattack. In particular, I want to clarify the relationship between concepts and the contingent features of speakers

on which they depend, so as to show that the relationship cannot be treated entirely as a matter of content. In my view, as I said, it is ultimately an empirical issue whether we should put ourselves in Cell Four. But to show that there is such an empirical issue, we need to do some philosophy.²

3. Two views of grounding

How should we think of the relationship between a declarative utterance and the “grounds” on which its key concepts depend? Take an example such as colour concepts. What is the relationship between an utterance such as “That’s red”, and what goes on in our visual system when we are presented with a ripe tomato? We can contrast two broad classes of answers to this question:

Keep the contingent grounds in the background

On this view, the proper place to mention the grounds is in a “use-condition”—a description of what typically or properly occurs in a speaker when such an utterance is made. The central idea is associationist: proper or normal uses “co-vary” with these conditions in the grounds. On this view, then, speakers acquire a habit of saying “That’s red” when certain circumstances obtain in themselves; but they don’t use this expression to say that these conditions obtain. Other familiar examples of this kind of view include Humean expressivist account of concepts such as causation, value, and probability, where the contingent grounds are psychological states of various kinds: habits or expectations, desires, and credences, respectively.

Of course, many aspects of such an account remain to be specified. What uses count, is the account descriptive or normative, and so on. However, for present purposes what matters is simply the contrast between this “use-based” way of theorising about the relationship between concepts and their contingent grounds, and an alternative approach.

Put the grounds in the foreground

On this alternative view, some mention of the grounds needs to be made in a fully explicit account of the content of the utterances in question—they are an aspect of what the speaker is talking about. In the colour case, for example, two representative versions of this view are (a) the subjectivist account, according to which “That’s red” reports the occurrence of a particular state in the speaker’s visual system; and (b) the dispositionalist

² To be more specific, I am interested in this possibility with respect to what we usually think of as the descriptive or representational uses of language. It isn’t controversial that language has multiple functions of other kinds, and some of these may well be hardware-dependent. (Competition for best example?)

account, according to which “That’s red” reports the existence of a disposition on the part of an external object to produce such states in the visual system. Either way, the grounds themselves—the states of the visual system—are regarded as an aspect of what the utterance is about.

Again, the Humean examples are helpful. Think of the familiar contrast between expressivism about value, probability or causation, on the one hand, and subjectivism (or “self-descriptivism”), on the other. The subjectivist says that utterances of these kinds are in part about the speaker’s psychological states.

Again, there are further issues that remain to be clarified. Of particular importance for our purposes is an issue concerning the notion of content invoked by such an account. Should we think of it “internally” or “externally”? If a concept is in part “about” its contingent grounds, is this to be thought of as a conceptual truth, accessible in principle to the average competent speaker? Or is it a kind of referential fact, accessible (at best) only a posteriori?

Eliding this last issue for the moment, the present relevance of the distinction between backgrounding and foregrounding views is that the latter count against Cell Four. If the grounds are “out in the open” in this way—part of what istalked about (in some sense) when concepts of the class in question are used—then their use need not be restricted to creatures who possess the relevant attributes themselves. We can’t use concepts whose use requires possession of an attribute we don’t have, but there seems no difficulty in talking about attributes which we don’t possess. (At rate, any remaining difficulties seem to be of one of two kinds: either merely epistemological, rather than something more basic; or dependent on the fact that the foregrounded reference to grounds is in part indexical, and therefore inaccessible to other speakers for that reason. I’ll set aside the latter possibility for the purposes of this talk.)

Thus foregrounding seems to challenge the possibility of hardware-dependence. It also challenges modularity, I think, for again, the supposedly distinct modules come to be treated as subspecialities of psychology, broadly construed. They simply describe different aspects of ourselves, or our relations to the world.

At first sight, the distinction between these two views may seem clear enough. But what precisely is the issue? The first thing to note is that the disagreement is not a simple choice between alternatives. In effect, the use-conditional account is the default position. Both sides agree that utterances of “That’s red” co-vary (more or less) with occurrences of certain grounding states (and more generally with various states of the environment). All

this “sideways-on” information about the use of the concepts in question is common ground, pretty much. (There may be some differences about how we should individuate concepts, for example.) There are lots of hard issues about which patterns are theoretically interesting, but these are issues for everybody. The real disagreement is about whether this common core provides the whole story about the relationship between the utterances and their contingent grounds; or whether, on the contrary, it is appropriate to say something extra—appropriate also to refer to the grounds in an account of the content, or truth conditions, of the utterances concerned (in either an internalist or an externalist sense)

This makes the dialectic non-symmetric. There is an important sense in which the supporter of the foregrounding view needs to make the running. Inter alia, she needs to convince us that the semantic notion (content, truth conditions, reference, or whatever) bears this kind of theoretical weight.

What are we asking when we ask whether an utterance has truth conditions, and if so what these truth conditions are; or when we ask what the content of assertion is? We tend to take for granted that the questions themselves are well-defined, even if the answers may be controversial in many cases. In my view, this confidence is misplaced. I’ve always been impressed by a remark of Dummett’s on this issue, in his paper “Truth” from 1959. “At one time”, says Dummett—referring to a long-gone philosophical era which must then have been six or seven years in the past— “It was usual to say that we do not call ethical statements ‘true’ or ‘false’, and from this many consequences for ethics were held to flow. But the question is not whether these words are in practice applied to ethical statements, but whether, if they were so applied, the point of doing so would be the same as the point of applying them to statements of other kinds” [TOE, p. 3].

In some ways, the question of what is at issue in such debates has not progressed greatly since 1959, I think. Without suggesting that its authors are especially at fault—in one respect, quite the contrary—I want to mention a recent contribution by Frank Jackson and Philip Pettit, a paper called “A Problem for Expressivism” (*Analysis*, 58(1998), 239–251). Apart from illustrating how it continues to be taken for granted that the issue is in good standing, this paper is useful for my purposes in a couple of ways. For one thing, Jackson and Pettit challenge—rightly, in my view—a popular opinion as to how the issue is to be decided. For another, they offer an alternative argument on the side of cognitivism. Although flawed, this argument fails in a useful way—a way which exposes one of the fundamental flaws of the foregrounding strategy.

4. A problem for expressivism?

Jackson and Pettit's target is expressivism (about ethical judgements, say). They invoke a Lockean principle to argue, in effect, that expressivism is always "trumped" by subjectivism—that is, by the view that moral claims report the very attitudes which the expressivist takes such claims to express. The argument goes like this:

- (i) Following Locke, we recognise that words are voluntary conventional signs; natural languages are not innate.
- (ii) To learn a convention, we must know what it is. If the expressivist is right, for example, we must know that the convention is that one say "X is good" when one approves of X (or "X is red" when one experiences a certain visual state when confronted by X, and so on). In other words, one must know that one's saying will be taken as an indication that one has the state concerned.
- (iii) Hence, in effect, such a conventional "saying" reports the presence of the corresponding attitude; it conveys this information to other members of one's speech community.

As Jackson and Pettit note, however, there is an obvious objection to this argument:

Expressivists, and philosophers in general, often rightly distinguish expressing what you believe from reporting what you believe. The sentence 'Snow is white', uttered in the right context, expresses your belief that snow is white, and is true iff snow is white. It does not report your belief. If you want to report your belief that snow is white, you need to use the sentence 'I believe that snow is white' (if you are speaking English), and this sentence is true iff you have the belief.

Expressivists often suggest that we can apply this distinction to other psychological states, including especially the 'ethical' attitudes, and that when we do, we get the account they need of the sense in which ethical sentences express attitudes. They observe that we can distinguish the doctrine that 'X is right' reports a certain pro-attitude to X, from the doctrine that it expresses that pro-attitude to X. The first view is subjectivism; the second, they claim, is expressivism. They may argue, then, that the availability of this distinction shows that there has to be something wrong with the argument from Locke. There has to be a sense of 'express' which ensures that 'X is right', and the like, express attitudes without reporting them and without having truth conditions.

The objection seems a strong one, and yet Jackson and Pettit think that it fails:

The trouble for expressivists is that, although there is an important difference between reporting and expressing a belief, it is plausibly a difference in what is reported. It is not a difference between reporting something and not reporting at all. When you express your belief that snow is white by producing, in the right context, the sentence ‘Snow is white’, you are not avoiding the business of reporting altogether. You are not reporting the fact that you believe that snow is white but you are reporting the content of that belief; you are reporting that snow is white. This is how the sentence gets to be true iff the belief is true.

If we take the distinction as drawn for beliefs, then, and apply it to attitudes, we do not get a result that can help expressivists. What we get is that ‘X is right’ expresses a certain pro-attitude iff ‘X is right’ reports the content of the attitude. And this is not at all what expressivists are after. First, it makes ‘X is right’ out to have truth conditions, namely, those of the content; and, second, it is very implausible in itself. The relevant content will be something like *that X happens* — for that is what we are favourably disposed towards, according to expressivists, when we assert that X is right — and that is very different from X being right (unfortunately).

However, this response relies on saddling the expressivist with a commitment which, while certainly implausible, is actually irrelevant to original objection. The original objection simply turns on the following point: Because we can express a belief without reporting that we hold that belief, the fact that words are “voluntary signs” for our psychological states cannot entail that in uttering those words we are saying that we hold the states in question. In other words, it does not entail that our own psychological state is part of the content of the utterances that express that state. So the general Lockean principle on which Jackson and Pettit rely is false—it must be, for it has this glaring counterexample. (Whether the fault is Locke’s or Jackson and Pettit’s need not concern us here.)

To rescue the argument, it is up to Jackson and Pettit to show us that the principle remains reliable in other cases—that there are no other relevant counterexamples. What they in fact show is merely that there are no counterexamples which work in exactly the same way as the one we have been given. This isn’t controversial, but nor is it relevant. Nobody thought that the counterexamples the expressivist needs would work in exactly the

same way. What maps over is simply the distinction between expressing and reporting an attitude, and Jackson and Pettit have done nothing to show that the expressivist is not entitled to that.

Here's an analogy. Suppose someone argues for theism from the premise that everything has a cause, suggesting that God is the cause of the universe as a whole. We point out that general principle is true, then God Himself has a cause—surely a reductio of the theist's position. Now imagine a theist who responds to this as follows: "I acknowledge (of course) that the principle that everything has a cause does not apply to God Himself—it doesn't apply to Divine things. But this is no use to my atheist opponent in establishing that worldly things might lack causes, for he doesn't want to say that worldly things might be Divine. (He doesn't want to allow that anything is Divine!) So my principle is safe in the cases that matter." This reply to the original objection works in the same way as Jackson and Pettit's reply, by attempting to saddle the opponent with the view that any other counterexample to the general principle at issue would have to have all the features of the one offered. In both cases, the onus is on the party seeking to rely on the generalisation to show us that it doesn't have other counterexamples. (As far as I can see, there is no way Jackson and Pettit could do that, short of coming up with an independent argument that expressivism is mistaken.)

5. Are there semantic facts?

Despite the untenability of Jackson and Pettit's main argument, the discussion is a useful one for my purposes. For one thing, it calls attention to the basic issue, which Jackson and Pettit do not themselves discuss: What precisely is at issue between them and their expressivist opponents. In what sense, if any, is it a determinate matter whether the expressivist or the subjectivist is right? As I noted earlier, the use-conditional approach is a fall back position, from which other people need to distinguish themselves. (This fall back position is not quite expressivism. Expressivists usually take for granted that there is a determinate and theoretically interesting sense in which the target sentences do not have truth conditions. The fall back position is neutral on this issue.)

In practice, one factor to which people often appeal is the ordinary use of the notions of truth and falsity to utterances of the class in question. For example, as Jackson and Pettit note, expressivists commonly argue that subjectivism has trouble making sense of the circumstances under which we are inclined to say that someone else's moral claim is "false". For example, if I say "Milosevic is a bad man", and you disagree, you don't seem to be challenging the claim that I disapprove of Milosevic. The argument doesn't seem to be

about my psychological state. However, as Jackson and Pettit rightly point out, these appeals to usage are very far from conclusive. There is a lot of scope for “explaining away” apparent anomalies.

Of course, Jackson and Pettit make this point against the background of the assumption, shared with their assumed opponents, that there is a fact of the matter whether the disputed utterances have truth conditions at all, and if so what these truth conditions are. But their own point tends to undermine this assumption. In effect, they argue that the ordinary use of the expressions “P is true” and “P is false” does not settle the issue of what the truth conditions of P are (or whether it has them). But given the standard equivalences—of “P” with “P is true”, and “Not-P” with “P is false”—it is very difficult to make this point at the semantically-ascended level without accepting it at the lower level as well. At that level—the level at which we say ‘P’, rather than “P is true”—it amounts to the claim that the use of the utterance “P” doesn’t make it a determinate matter what an utterance’s truth conditions are (or that it has truth conditions at all). But is usage doesn’t fix it, what else could do so? Why should we suppose that there is a matter of fact to be fixed?

6. The inevitability of backgrounding

I will come back to these issues below, for they connect with the issue of referential externalism, which I have so far deferred. For the moment, let’s consider a counterargument, which will take us back to my main theme. It might be objected that if we didn’t know what the issue was about possession of truth conditions, we couldn’t have been so sure that ordinary assertions are not about the beliefs from which they stem; that the utterance “Snow is white” does not have the content “I believe that snow is white.” Well, how did we know this? Did we simply rely on our intuitions that “snow is white” is not a self-description, not about our own psychological state? If so, then perhaps Jackson and Pettit should not have taken for granted that “Snow is white” is not (at least in part) a self-description. After all, we seem to have the same intuitions about the cases they do want us to treat as self-descriptions. (Imagine an expressivist about “white”, for example.) (More to the point, why were Jackson and Pettit entitled to help themselves to the assumption that truth conditions of ‘snow is white’ are that snow is white? Why not take this to be another case in which surface usage is misleading. Perhaps the ‘real’ truth conditions of ‘snow is white’ are conjunctive: ‘snow is white and I believe that snow is white’, for example.)

In fact, we can do better than simply appealing to intuition. We can point out that the move to interpret all assertions in this self-referential way leads to a disastrous regress.

The same general principle which tells us that “Snow is white” has content “I believe that snow is white” would tell us that the latter statement has content “I believe that I believe that snow is white”; and so on. Of course, none of the resulting iterating series of statements is unacceptable in itself. What is unacceptable is just that we should be forced to keep on going. If a principle implies that we do keep on going—even if only implicitly, in the sense that we commit ourselves implicitly to all these things, in asserting “Snow is white”—then that principle must be false.

This logical point shows us that something must be backgrounded—the relation between belief and assertion cannot be such that beliefs go into the foreground, into the content of the assertions concerned. This is helpful, from my point of view, but it doesn’t establish that the same thing might be true in a more local way—that there might be particular backgrounding relationships between the use of particular concepts and the grounds on which they depend. To get to this more local conclusion, we need to appeal to considerations which are more empirical than logical. One key point is the familiar observation that language acquisition cannot be entirely explicit—we can’t learn everything by learning definitions. Some concepts have to be picked up on the fly—we simply have to acquire the right habits. We have to learn to align our linguistic dispositions with other aspects of our psychology, without ever being in a position to say or conceptualise the thought that that is what we are doing. In these basic cases, then, the relation between psychological grounds and the concepts themselves seems inevitably a backgrounding one. We do not acquire a concept which in part concerns our psychological state. To that extent, then, the Jackson and Pettit version of Locke must be mistaken.

But now externalism intrudes again, objecting that what we are aware of or capable of conceptualising is irrelevant. Our concepts may in part “pick out” something psychological, even if we ourselves are not aware that this is the case. Again, then, this is an appeal to the externalist strategy, which I have deferred several times. In the next section, I want to argue that externalism cannot be properly motivated, once linguistic imprisonment is disallowed. If externalism can be dismissed in this way, we will have the kind of conclusion I want. There must be local backgrounding in language, at the level of individual concepts. Which concepts? I don’t know, but neither does anyone else. Are they hardware dependent, and/or modular? Again, we don’t know, but my point is that these are empirical matters. My main purpose here has simply been to defend this space of empirical possibilities against a philosophical move which would close it off. Externalism is the remaining threat.

7. Externalism regained?

The basic externalist thought is something like this. Differently-imprisoned speakers might nevertheless be referring to the same things. Differently-grounded concepts may nevertheless 'hook onto' the same world.

From what stance does the externalist speak? Not from outside language, on pain of violating imprisonment. So to the extent that we can be externalists, in the light of linguistic imprisonment, we must be able to do it from where we stand, within the systems of concepts we possess. Externalism must speak "from within" about the relations between language and the world.

Still, what is the problem with this? Isn't the use perspective already externalist in this sense? After all, it is about covariance patterns between linguistic items and non-linguistic items. If the externalist finds reference relations (and the like) in those patterns, the use-theorist can hardly object. True, the resulting story is in one sense already a part of the use-theory. But mightn't this be like the sense in which chemistry is already a part of physics? In one sense it may be true, and yet chemistry is worth formulating independently.

But let's think more carefully about this project of finding the semantic relations within the covariance patterns, broadly construed. First, what is at stake when we make our choice—what constraints are there on which pattern we should count as the reference pattern, for example? Obviously, it is not supposed to be a stipulative exercise. We can't simply pick out part of the pattern and call it "reference"—the project begins with the assumption that there is a right answer (near enough). And the constraint is surely the one imposed by the ordinary use of the term "reference" (or by some properly specified theoretical use, if the notion we have in mind is not thought to be the folk notion). The constraint comes to this: the reference relation is whatever these uses of the term "reference" actually refer to. (This could be put in terms of a Ramsey sentence.)

Now the problem is obvious, isn't it? In taking this as a constraint, we assume the very claim at issue, namely that there are determinate semantic facts. Without this assumption, the ordinary use of the semantic terms cannot be taken to "pick out" certain features of patterns of usage (or anything else, for that matter), except in the trivial ways specifiable from within: the term "reference relation" refers to the reference relation, and such like. In other words, the project assumes with respect to the semantic terms themselves the very issue which is contentious with respect to terms in general, namely,

that there is some theoretically interesting sense in which reference and truth conditions are determinate.

In fact, even if we ignored the circularity at this point—i.e. its role in the description of the motivation of the externalist task—it would emerge again, at the level of any proposed solution. For in effect, the project is to decide what natural pattern the referential concept of “hooking on to” actually hooks on to—and this is simply an equation with too many variables. (This is essentially Putnam’s argument, I think, though perhaps in a different dialectical context.) So the externalist strategy is doubly circular, and provides no well-grounded escape from the threat of linguistic imprisonment.

8. Conclusion

I have argued that there must be local backgrounding in language, at the level of individual concepts. We don’t know which concepts they are, except in generic terms: they are those that do not have an available analysis within the framework as it stands. It is essentially an empirical matter which concepts these are, even if one we can investigate to some extent by introspection. Likewise, it is an empirical matter whether they are hardware dependent or modular. As I said, my main aim here has been to defend this space of empirical possibilities against a philosophical move which would close it off.

All the same, I can’t resist a bit of armchair biology. I think it is very likely that many concepts will turn out to be both hardware-dependent and modular, in interesting ways. Hardware-dependence seems by far the cheapest solution to the kinds of problems which confront emerging language users. The essential point is that there’s no need to represent what everybody has in common. It is wasteful and counterproductive to represent what doesn’t change. (If we all had rose-tinted corneas, for example, it would be pointless to qualify our colour judgements by saying that they were made from that rosy perspective—far cheaper to divide through by what we have in common.) So hardware-based commonalities are likely to be backgrounded, in general. We won’t develop conceptual resources where they are not needed.

Moreover, given the wide variety of very different kinds of hardware commonalities, this source of hardware-dependence seems likely also to give rise to inhomogeneity, or modularity. There are many different ways in which we are all alike: in terms of each of our sensory modalities, our temporal experience, the fact that we are agents, and so on. If we are imagining a space of possible speakers, each of these aspects seems to be an independent dimension of variability (more or less). So if each is the background to some group of concepts, the resulting language is thereby modular. Roughly,

each group of concepts has a function which reflects the peculiar needs of creatures whose hardware places them in a particular kind of relation to their environment.

Thus I think that language is very likely to turn out to exhibit the kind of perspectival character associated with the fourth cell in the table mentioned earlier. If so, then linguistic imprisonment does represent a real constraint, in the sense described at the beginning. It implies that our conceptual viewpoint is inevitably a product of our physical circumstances. Had our circumstances been different, we would not have occupied this viewpoint. Our viewpoint is inevitably “situated” in ways which depend on our physical circumstances. In this sense, then, there is no such thing as a view from nowhere—a viewpoint which is “unsituated” in these ways. In particular, we cannot hope to achieve such a viewpoint by simply being explicit about the subjectivities—by trying to render them all in conceptual content.

This is not to deny that we may be able to make changes of degree, by restricting ourselves for certain purposes to concepts which do not exhibit particular kinds of dependence. It may be sensible to keep humour out of metaphysics, for example. More seriously, as I’ve argued elsewhere, it may be sensible to keep our familiar embedded temporal perspective out of discussion of issues about the physics of time asymmetry—to try to address these issues from what I called a view from nowhen. One very interesting question is how far this project of “sterilising” language for particular purposes can and should be taken. There is a tendency to think that if we followed this path to its endpoint, we would reach a genuinely non-perspectival language. This tendency calls for two loud cautions, I think.

First, if our goal is an ideal scientific language, then it is far from clear that less perspectivalism is always better. Suppose it turns out that our causal and modal concepts are grounded in an Humean manner on contingent features of our circumstances: the fact that we are agents and deliberators, perhaps. Does this mean that science would be better off without these concepts? An alternative possibility is that science is, *inter alia*, the kind of activity which is only possible from the perspective of agents and deliberators. If so, then surely we should embrace the perspectival character of scientific language, rather than trying to eliminate it.

The second caution is that if the argument of this paper is correct, there is an important sense in which this idealising project is misconceived. Since there is no such thing as a language without backgrounding, there is no such ideal language to be reached. We can’t reach a view from nowhere by subtraction, as it were—by trying to leave out just

the perspectival bits. In the Sydney real estate terms which inspired my title, the analogy goes like this. In general, one improves one's harbour view by moving closer, or by moving higher. But neither strategy leads in the limit to the harbour view to trump all harbour views. In the limit, in both cases, we end up with no view at all.

Should we find it disturbing that linguistic imprisonment confines us in this way—that it traps us in our physical embodiment, when we try to get on with the mundane but important work of saying how things are? It may be a matter of temperament. It doesn't bother me, but perhaps I'm abnormal. In contemporary philosophy the people most likely to be bothered by it are those who see it as a threat to their preferred form of realism. For such people, the point to be emphasised is that the conclusion itself results from taking seriously a mundane naturalistic view of ourselves. Given linguistic imprisonment, it turns out that naturalism and perspectivalism go hand in hand.