



Huw Price

“On the Road with Rorty, Davidson, and Ramberg”

In: Yvonne Huetter–Almerigi and Robert Sinclair (Eds.) (2026).
Pragmatism, Metaphysics and Method—Essays for Bjørn Ramberg
(pp. 135–172). Nordic Studies in Pragmatism 5. Helsinki: Nordic
Pragmatism Network.

ISSN-L 1799–3954

ISSN 1799–3954

ISBN 978–952–67497–4–7

Copyright © 2026 The Authors and the Nordic Pragmatism Network.

This work is licensed under a Creative Commons
Attribution-NonCommercial 3.0 Unported License.

 For more information, see
<http://creativecommons.org/licenses/by-nc/3.0/>



Nordic Pragmatism Network,
Helsinki 2026

www.nordprag.org

On the Road with Rorty, Davidson, and Ramberg

Huw Price
Trinity College, Cambridge

1. Encounters with Rorty

In a recent piece, Bjørn Ramberg remembers Rorty as a generous philosophical mentor (Ramberg, 2022). I was envious. Rorty would have been high on my wish list, if I'd been the mentorable type.

As it is, I met Rorty only three times. The last was just a few weeks before his death, in 2007. Brandom repeated his Locke Lectures in Prague, over two days, and I was a commentator. Brandom told me how moved he was that Rorty, gravely ill, had insisted on being there.

The first two meetings were in Australia, in July 1999. Rorty was visiting ANU's Humanities Research Centre. We met first in Melbourne, at the annual conference of the Australasian Association of Philosophy. I gave the Presidential Address that year, and Rorty an evening keynote, 'Some Doubts about the Idea of "Philosophical Method"'. I can't recall the precise doubts in question, but I remember my sense of the Q&A—Rorty handling a large and rather hostile audience like a shambling grandmaster, playing several games simultaneously.

The title of my lecture, 'Location, Location, Location', was a nod to the location of the conference itself. In Melbourne it is said, not always in jest, that we Sydneysiders care about nothing but real estate and harbour views. I was indulging this stereotype, for my local hosts. The play must have seemed parochial as well as feeble to Rorty, but he was generous about the talk.

Despite Rorty's seal of approval, the lecture never made it into print. A few years ago I put a copy online (Price, 1999; hereafter LLL), to cite

in another piece. When Yvonne Huetter-Almerigi and Robert Sinclair approached me about contributing to the present volume, we soon triangulated, as it were, on the proposal that LLL be the basis of it. Huetter-Almerigi had shown me work of hers, relating my views to Davidson's famous triangulation argument. That got me thinking about the relevance of Davidson's externalism to the conclusions of LLL. It became clear that those issues had links not only to LLL, but also to Ramberg's work on Davidson and Rorty.

Some of these connections also link to my second encounter with Rorty in 1999. A few days after the Melbourne conference, ANU hosted a workshop with Rorty. I presented the original version of 'Truth as Convenient Friction' (Price, 2003; hereafter TCF). This piece takes up Rorty's objections to Crispin Wright's views on the normativity of truth (Rorty, 1995). Rorty begins that paper with a pragmatist maxim:

Pragmatists think that if something makes no difference to practice, it should make no difference to philosophy. This conviction makes them suspicious of the distinction between justification and truth, for that distinction makes no difference to my decisions about what to do. Rorty, 1995, 19

He argues that Wright's view falls foul of this principle. In response, drawing on proposals in earlier work (Price, 1983; 1988; 1998a), I argued that a norm of truth distinct from justification actually makes a huge behavioural difference:

What [...] [such a] norm provides [...] is the automatic and quite unconscious sense of engagement in common purpose that distinguishes assertoric dialogue from a mere roll call of individual opinion. Truth is the grit that makes our individual opinions engage with one another. Truth puts the cogs in cognition, at least in its public manifestations. TCF, 165

About this paper, Rorty gave me comments at more than one point. An exchange based on those remarks was published as (Rorty and Price, 2010). Rorty invokes Davidson against a key move in my argument; and this point, too, has links to Ramberg's work on Davidson and Rorty.

In the light of these various connections, the present piece is constructed like this. Its two supporting axles are LLL and TCF, the two pieces I first presented in 1999, with Rorty in the room. Riding on these wheels is

a cluster of issues from the work of Davidson, Rorty, and Ramberg. Since most of these concerns turn on the sense in which language is necessarily a *social* activity, it seems appropriate to think of the vehicle as a Volkswagen. (If we make it a Beetle, we can put Wittgenstein in the driver's seat.)

I'll begin with the more obscure of the two axles, namely LLL. Section 2 reproduces LLL as originally written, with a few tweaks to references and footnotes. Section 3 is then in three parts. § 3.1 begins with an objection to my conclusions in LLL, based on some of Davidson's remarks about externalism and triangulation. I don't claim that Davidson would endorse this objection. On the contrary, as we'll see in § 3.2, it is clear that he would not—and partly for reasons that Ramberg emphasises, in criticising Rorty's interpretation of Davidson. Those reasons link in turn to Rorty's invocation of Davidson, in response to TCF. As I'll explain, I think this invocation backfires for Rorty, supporting the main claim of my paper.

I realised recently that in responding to Rorty on this point, I could have appealed to Bernard Williams's (2002) characterization of 'imaginary genealogy'. For this, Williams's model is Edward Craig's genealogy of the concept of knowledge (Craig, 1990). Yet, citing Davidson, Williams declines to seek what TCF claims to offer, a genealogy of truth itself. In § 3.3 I'll argue that TCF thus has lessons for Williams and Davidson, as well as for Rorty. By this roundabout route, our bespoke VW reaches a familiar destination. I conclude that all of its famous passengers have something to learn from Cambridge pragmatism (Misak, 2016; Price, 2017).

2. Location, Location, Location

2.1 Linguistic imprisonment

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.

It might be argued, as in effect by Davidson (1974), 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 non-threatening 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 (see § 2.7 and § 3.1), 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.

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.)¹ 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.

¹ In recent work (Price, 2017; 2023b) I use the term 'pragmatic grounds'.

2.2 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 (e.g., Price, 1997), 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.²

Thus the two distinctions are independent, giving us a four-fold table of possibilities (see Table 1, next page). 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

² This distinction might become blurred if we softened the notion of hardware, and recognised 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.

any modularity is “non-vicious”—in principle, any speaker *could* run any module, even if not all at once. But the combination of multi-functionalism and hardware-dependence seems to leave us especially trapped. That’s what makes the fourth option particularly interesting.

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

In my view, it is an empirical issue whether our own constitution and the nature of language puts us in Cell 4. 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 4. It argues that it is always possible to interpret the *content* of concepts in a way which renders them non-hardware-dependent and uni-functional. The other counterattack claims to mitigate the philosophical consequences which would follow if we were in Cell 4. 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 above (§2.1). 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 (see §2.7 and §3.1), 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 4. But to show that there is such an empirical issue, we need to do some philosophy.³

³ 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

2.3 Two views of grounding⁴

How should we think of the relationship between a declarative utterance and the contingent 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 (see Price, 1991 for further discussion).

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 account, according to which “That’s red” reports the existence of a

language has multiple functions of other kinds, and some of these may well be hardware-dependent.

⁴ In 1999 there was less risk than there is today that this might be taken to refer to what is now known as metaphysical grounding. As the context makes clear, I mean something very different: a relation between utterances, or linguistic items of some other kind, and contingent features of speakers on which their use depends.

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 now, the relevance of the distinction between backgrounding and foregrounding views is that the latter count against Cell 4. If the grounds are “out in the open” in this way—part of what is *talked 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 piece.)

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 sub specialties 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 environ-

ment). 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, as to a long-gone philosophical era, to a period 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 [...]

Dummett, 1959, 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 piece by Frank Jackson and Philip Pettit (Jackson and Pettit, 1998). 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.

2.4 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.

Jackson and Pettit, 1998, 244–45

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).

Jackson and Pettit, 1998, 245

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 wordly things might lack causes, for he doesn’t want to say that wordly 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.)

2.5 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 if usage doesn't fix it, what else could do so? Why should we suppose that there is a matter of fact to be fixed?

2.6 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.)⁵

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.

⁵ More to the point, why were Jackson and Pettit entitled to help themselves to the assumption that the 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.

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.

2.7 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 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. 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.⁶) So the externalist strategy is doubly circular, and provides no well-grounded escape from the threat of linguistic imprisonment.

⁶ I discuss this comparison in (Price, 1998b), and the more general point in (Price, 2004).

2.8 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 judgments 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 (Price, 1996). 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 we should embrace the perspectival character of scientific language, rather than trying to eliminate it.⁷

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.

⁷ For more on these themes, see (Price, 2007; 2023b).

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.⁸ Given linguistic imprisonment, it turns out that naturalism and perspectivalism go hand in hand.

3. Externalism, community, and objectivity: three points on the Davidsonian triangle

I now want to make three points, related to Davidson's triangulation argument and to my two pieces (LLL and TCF) from 1999. I'll begin by asking whether Davidson's version of externalism is a threat to the thesis of LLL. The answer turns out to be no, but there is enough in Davidson's formulation of triangulation to suggest a superficial case for the other side. I want to explain how LLL can respond, even though it is not Davidson's actual view.

Second, I'll note that Davidson's own objections to the version of externalism in question take us to aspects of his view that Ramberg has emphasised, against Rorty. It also takes us to Rorty's invocation of Davidson, against a central move in my TCF. As I'll explain, I think the point backfires for Rorty. I'll back up this response with some observations from Bernard Williams (2002) about 'imaginary genealogy'.

Williams himself declines to seek such a genealogy for truth itself, partly for reasons he finds in Davidson. For my third point, I'll argue that the genealogical approach to truth I advocate in TCF has much to offer to Davidsonians. It offers significant insights about truth, without violating Davidson's prohibition on analysis or reduction.

3.1 Is Davidson's triangulated externalism a threat to LLL?

For the first point, it will be helpful to have on the table a distinction between two distinct notions that tend to be lumped together, under 'representation' and related headings. I have proposed that we distinguish

⁸ A *subject* naturalist view, as I came to call it (Price, 2004).

two ‘nodes’ in work on representation in philosophy, cognitive science, and related fields (see Price, 2011a; 2015, Price *et al* 2013).⁹ Each of these nodes is itself a cluster of notions, in the sense that the defining features of the node can themselves be developed in a number of different ways. Bracketing the latter diversity, my proposal is that we see the nodes themselves as distinct notions, rather than different aspects of the same single concept of representation.

The first node, *e-representation*, involves a world-tracking conception of representation. This is familiar in biological cases, often in association with the idea that the function of evolved representations is to co-vary with some (typically) *external* environmental condition. There are also non-biological examples, of course: the barometer reading and air pressure, and so on. What unites such cases is that some feature of the representing system is (in some sense) intended to vary in parallel with some feature of its environment, exploiting some causal mechanism.

E-representation thus emphasises the *system–world* links. The second node, *i-representation*, gives priority to *internal* connections between one representation (such as a word) and another. By this criterion, a token counts as a representation in virtue of its position in some sort of cognitive, inferential or functional architecture—its links, within a network, to other items of the same general kind. The notion is flexible enough to accommodate several different kinds of network: causal–functional, inferential, even computational. It might also be social, not internal to any single speaker. The sense of ‘internal’ that matters involves looking to the network, not to the world. The important thing is that the notion be divorced from any external notion of representation, thought of as a system–world relation.

These two notions of representation may have different uses, for various theoretical purposes. By the time we get to human language, from simpler forms of biological representation, in either sense, there may be no useful external notion of *a semantic kind*—no useful general notion of a relationship that words and sentences bear to the external world, that we might identify with truth and reference. The impression that there are such external relations can be regarded as a kind of trick of language—e.g., a misunderstanding of disquotational platitudes. We can hold this without rejecting the internal notion—without thinking that there is no interesting sense in which mental and linguistic representations can be

⁹ (Gibbard, 2015) and my response in (Price, 2023a) may also be helpful.

characterised and identified in terms of their roles in networks of various kinds.

Some critics are unhappy with my ecumenical use of the label ‘representation’, in drawing this distinction. They agree with me that there is such a distinction, but insist that only one of my two categories really deserves to be called representation. However, such critics disagree amongst themselves: some say that e-representation is the real thing, others i-representation. A third group responds that any genuine representation would need both faces, agrees that these are different things, and concludes that there are no representations at all.¹⁰ I take all of these reactions to confirm my main point. Terms such as ‘representation’ are used in a variety of ways in philosophy and related fields, and we need to be careful, lest we talk past each other.

As noted, e-representation is often a causally-grounded notion. Organisms, human or otherwise, use such representations to coordinate with, or *track*, relevant features of their (usually) external environments. And causation makes this possible. When things are working as they should, e-representations are *effects* of the environmental conditions they have evolved to track.¹¹

Such a view is externalist in an obvious sense. It invites us to identify the content of internal states with the external circumstances they function to track. Applied to linguistic e-representations, then, it seems to challenge the ‘located’ view of language described above. Doesn’t causality reach outside our linguistic prisons?

It may seem easy to read this as a Davidsonian objection. Davidson certainly says that he is an externalist, and his crucial notion of *triangulation* invokes easy talk of causation and the world.

[E]ach of two people is reacting differentially to sensory stimuli streaming in from a certain direction. Projecting the incoming lines outward, the common cause is at their intersection. If the two people now note each other’s reactions (in the case of language, verbal reactions), each can correlate these observed reactions with his or her stimuli from the world. A common cause has been determined. The triangle which gives content to thought and speech is complete. But it takes two to triangulate.

Davidson, 1991, 213, underlining added

¹⁰ It would be easy to read Davidson (1990) as an example.

¹¹ Millikan (1989; 2009) provides a famous example of such a view.

The last sentence here alludes to the most important reason why Davidson is *not* an e-representational externalist, in my sense. We'll come to that, but first to some more direct objections to any attempt to use triangulation as the basis of an externalist challenge to the view of LLL.

Davidson's triangulating creatures are similarly 'located', in my sense, in that they have similar sensory and perceptual responses to their environment. Think of such responses, on the part of any single individual, as the manifestation of dispositions to respond to the passing scene in certain ways. Triangulation requires that two individuals have similar dispositions. Crudely, then, what they are triangulating *on* is something like the world *as perceived by a creature with such dispositions*. The dispositions in question count as a location, or contingent ground, in my sense. We don't put them into the foreground, into the account of content given to us by triangulation, but there is no pretending that they are not there, unpinning the story.

Davidson is himself explicit about the need for similarity, on the part of the creatures doing the triangulation:

[F]or starters we learn words for what we preverbally discriminated; we learn from creatures very like ourselves, and could not learn from creatures very unlike ourselves.

Davidson, 2001d, 15

This is especially clear when we get to values. For all his talk of causation, Davidson's world is not the e-representationalist's natural world.

Gauging the thoughts of others requires that I live in the same world with them, sharing many reactions to its major features, including its values.

Davidson, 1991, 220

If we wanted to reinforce this point further, then, picking up on remarks in § 2.8, we could appeal to the perspectival character of the notion of causation itself—a notion that plays such a fundamental role in Davidson's picture (and about which, as about truth, he is anti-reductionist; see, e.g., Davidson, 1996, 20–21).

In conclusion, Davidson himself is not an e-representational externalist. As I have said, we'll come to an additional reason for this assessment in a moment. But the points here show that an e-representational caricature of

Davidson, put together from the most suitable fragments of his writings on triangulation, would not be much of a threat to the conclusions of LLL.¹²

3.2 Davidson's social externalism and Rorty's objection to TCF

We don't need to dig very deep to find Davidson's own reasons to dismiss an e-representational interpretation of his externalism. As he puts it himself:

There are two major species of externalism: *social externalism*, which maintains that the contents of our thoughts depend, in one way or another, on interaction with other thinkers; and *perceptual externalism*, which holds that there is a necessary connection between the contents of certain thoughts and the features of the world that makes them true. Davidson, 2001, 4

Davidson's own externalism is the social variety, as he makes clear.

Even Rorty doesn't always understand how deeply this commitment runs in Davidson, as Ramberg points out. Ramberg comments on Rorty's interpretation of the following remark from Davidson's 'Three Varieties of Knowledge':

A community of minds is the basis of knowledge; it provides the measure of all things. It makes no sense to question the adequacy of this measure, or to seek a more ultimate standard. Davidson, 1991, 218

Discussing Rorty's interpretation of this remark, and its surrounding context, Ramberg concludes like this.

¹² Rorty himself makes a similar point about Davidson in *Philosophy and the Mirror of Nature*.

Correspondence, for Davidson, is a relation which has no ontological preferences—it can tie any sort of word to any sort of thing. This neutrality is an expression of the fact that, in a Davidsonian view, nature has no preferred way of being represented, and thus no interest in a canonical notation. Nor can nature be corresponded to better or worse, save in the simple sense that we can have more or fewer true beliefs. Rorty, 1979, 300

Functional pluralists may raise an eyebrow or two at the use of 'nature' here (as we would to 'the world' in the same position). Otherwise, Rorty's point meshes well with the one I am making above.

This [...] suggests to me that Rorty gets the emphasis wrong when he glosses Davidson's point that a "community of minds is the basis of knowledge" as the claim that "knowing our way around the community of beliefs is the *basis* of the rest of our knowledge." [Rorty, 1999, 587] And it is because he gets the emphasis wrong in this way that he ends up with a construal of Davidson's key point in "Three Varieties of Knowledge" that is in tension with other pragmatic Davidsonian theses. In fact, Davidson is not holding one particular kind of knowledge up as the basis for other kinds; as Rorty sees and argues, the thought that any kind of knowledge is in that sense *basic* is quite at odds with the general thrust of Davidson's paper—indeed with his antirepresentationalism in general. The basis of knowledge, any form of knowledge, whether of self, others, or the shared world, is not a community of *minds*, in the sense of mutual knowledge of neighboring belief-systems, as Rorty here takes it to be. Rather, it is a *community* of minds; that is, a plurality of creatures engaged in the project of describing their world and interpreting each other's descriptions of it. This suggestion entails no priority of one kind of knowledge of the sort that Rorty's reading suggests, but emphasizes the intersubjective nature of *all* knowledge.

Ramberg, 2000, 361–62

There is some irony, for me, in Rorty's failure to appreciate the depth of Davidson's reliance on community. For it is just this aspect of Davidson's views that Rorty invoked in his objections to TCF. There, as in an earlier piece (Price, 1998a), I proceed by distinguishing the conversational norm I take to be associated with truth and falsity from two other norms: *sincerity* and *warranted assertibility*. Truth is thus my 'third norm', and the focus of the paper is on what this third norm adds to our conversational practice—on the question how things would be different if we didn't have it.

As an aid to answering this question, I imagine a speech community who don't have the third norm, but who nevertheless use speech acts to give voice to their opinions. I call such speech acts 'merely opinionated assertions', or MOAs, and imagine them used by a community called the Mo'ans. As I point out, it is easy to get to something like MOAs in our own

linguistic practices, by using devices to cancel the third norm: ‘My own opinion is that P’; ‘Mine is that not-P.’¹³

I was aware, of course, that it is doubtful whether the Mo’ans are a realistic possibility. But I felt that these doubts did not invalidate their use as a kind of thought experiment. Rorty focussed on this point, in the comments later published in (Rorty and Price, 2010).

Price asks us to imagine a community in which there are no attempts at intersubjective justification, but in which its members nevertheless express “the kind of behavioral dispositions which we would characterize as beliefs [...] by means of a speech act we might call the *merely-opinionated assertion* (MOA, for short)” [Price, 2003, 172]. He admits that one might doubt the possibility of such a community: perhaps, he says, “a truth-like norm is essential to any practice which deserves to be called linguistic” [Price, 2003, 168, n. 8]. But he thinks this possibility irrelevant to his thought experiment.

It seems relevant to me. I doubt that we can tell a plausible story about a Mo’an community. In particular, I do not see why a radical interpreter would construe as *assertions* the noises made by organisms that never attempt to correct one another’s behavioral dispositions—never try to get others to make the same noise they do. I would advance arguments familiar from Wittgenstein, Davidson, and Brandom to urge that there must be social cooperation on projects of shared interest before language can get very far off the ground. One cannot justify by own lights if one does not know what it is to justify by the lights of others. Price’s “chatter of disengaged monologues” [Price, 2003, 166] is possible only as an enclave within a culture in which there is lots of engaged dialogue.

Rorty and Price, 2010, 255

To me, this response seemed to count in my favour. In TCF I had argued, against Rorty’s objection to Wright, that the third norm makes a very big behavioural difference. As I put it: ‘It gives disagreement its immediate normative character, a character on which dialogue depends, and a character which no lesser norm could provide’ (TCF, 164). Here, Rorty seemed

¹³ In Wright’s (2021) terms, then, the Mo’ans are speakers who don’t, or no longer, have the device of putting their opinions into a public realm.

to be conceding that point, in objecting that the Mo'an's washed-out substitute for dialogue was simply impossible. Responding to Rorty, I put the point like this.

For my purposes, what mattered about the Mo'an's was that by seeing what their linguistic practice would lack, we see what truth adds to our own. What's missing for the Mo'an's—what the third norm provides for us—is (as I put it) “the automatic and quite unconscious sense of engagement in common purpose that distinguishes assertoric dialogue from a mere roll call of individual opinion.” Let's agree, with Rorty, that when we consider the Mo'an's in the light of “arguments familiar from Wittgenstein, Davidson, and Brandom,” we realize that there can be no such community. Removing that sense of engagement amounts to removing anything that might count as an assertion, or indeed as an expression of opinion, in the full-blown sense of the term. This is no reason to forget the lesson we learnt by trying to imagine the Mo'an's—on the contrary, as in many cases, the point of the thought experiment lies *precisely* in the fact that it leads us, in thought, to an impossible destination. (The lesson lies in the nature of the impossibility.)

Rorty and Price, 2010, 257

This reply to Rorty was written around 2008, in response to comments that Rorty had provided in 2005. I now realise that had I read Bernard Williams's *Truth and Truthfulness* (Williams, 2002) at that point, I could have put my response in his terms. As its subtitle says, Williams's book is an essay in philosophical genealogy. Williams offers a genealogy of what he calls *truthfulness*—roughly, the fact that humans in normal communities have a tendency to *value* truth. This is not a genealogy of *truth itself*, and we will come to that point shortly. Williams invokes Davidson, in setting aside such a project.

For the moment, what's relevant is the way in which Williams characterises his approach. He draws on that of his former Cambridge colleague, Edward Craig, in *Knowledge and the State of Nature* (Craig, 1990). Under the heading ‘How Can Fictions Help?’, Williams describes Craig's account like this.

Craig's example, like my own State of Nature story, is an example of what I shall call an “imaginary genealogy”—“imaginary,”

because [...] there are also historically true genealogies. Imaginary genealogies typically suggest that a phenomenon can usefully be treated as functional which is not obviously so. Moreover, they resemble a larger class of explanations (including those given by natural selection theory) in explaining the functional in terms of the non-functional or, perhaps, in terms of the more primitively functional. The power of imaginary genealogies lies in introducing the idea of function where you would not necessarily expect it, and explaining in more primitive terms what the function is. Williams, 2002, 32

It is this characterisation that I could have cited in my response to Rorty.

What does Williams mean by ‘function’ here? He introduces the term for the first time in the preceding paragraph, at the beginning of his discussion of Craig’s example.

[Craig’s] State of Nature gives an explanation, but what it explains, the concept of knowledge, does not look as though it had been designed. On the contrary: before the story, one may well never have asked what the function of the concept is, and that is part of the point. Craig’s story answers, in its fictional terms, the question “Why should we have a concept such as the concept of knowledge?” and, by answering it, suggests the question itself. What the question introduces is the notion of function, and that step itself does some of the work. If one sees the concept of knowledge as having a function—in particular, a function in relation to very basic needs—this in itself helps one to see why it has the features it has, and can discourage one from less fruitful approaches. Williams, 2002, 31–32

To my eyes, this gloss looks very congenial. It is what I had in mind in *Facts and the Function of Truth* (Price, 1988) and other pieces. But, unlike Williams and Davidson, I wanted to take this stance with respect to truth itself. (I was interested in the ‘point’ of truth, in the sense I took from the remark of Dummett I quoted in § 2.3 above.) The last issue I want to put on the table is why Davidson never seems attracted to such a genealogical approach. So far as I can see, he never has genealogy on his radar, whether about truth or any of the other notions he regards as similarly basic. Let’s approach this question via Williams, who certainly does have genealogy in view—very much so—but appeals to Davidson in resisting it in the case of truth.

3.3 Williams, Davidson, and the Function of Truth

Williams emphasises that he is seeking to explain *human concerns with truth*, not truth per se. He makes dismissive remarks about the possibility of making truth itself the focus.

The inquiry is, then, into human concerns with the truth. A basic form of that concern lies in the virtues of truth. [...] One thing I shall not consider [...] is the history of *the concept of truth*, because I do not believe that there is any such history. The concept of truth itself—that is to say, the quite basic role that truth plays in relation to language, meaning, and belief—is not culturally various, but always and everywhere the same. [...] There are indeed scholarly books that describe themselves as histories of the concept of truth, but they typically describe conceptions, varying over time, of belief-formation, or of knowledge, or of the metaphysical structure of the world. Often, they are histories of philosophical theories of the truth. [...] The present point is that philosophical theories of truth [...] quite certainly have a history, whereas the concept of truth itself does not.

Williams, 2002, 61–62

These are surprising claims for a genealogist. There's a trivial sense in which the concept of truth has a history, because, like all our concepts, our pre-linguistic ancestors did not possess it. In the light of this, the claim that truth is 'always and everywhere the same' seems cavalier. How could we know, if we're ignoring that part of the history?

Why does Williams take such an uncompromising stance on these issues? He says that he is following Davidson's lead.

What about truth itself? If we are going to say that beliefs and assertions [...] aim to be true, or, as the State of Nature story has assumed, that it is a good idea [...] that they should be true, should we not say something about what it is for them to be true?

We should say something, but not very much. In particular, we should resist any demand for a *definition* of truth, principally because truth belongs to a ramifying set of connected notions, such as meaning, reference, belief, and so on, and we are better employed in exploring the relations between these notions than in trying to treat one or some of them as the basis

of the others. It is also true that if any of these notions has a claim to be more basic and perspicuous than the others, it is likely to be truth itself. As Davidson has put it:

[W]e cannot hope to underpin [the concept of truth] with something more transparent or easier to grasp. Truth is, as G. E. Moore, Bertrand Russell and Gottlob Frege maintained, and Alfred Tarski proved, an indefinable concept. This does not mean we can say nothing revealing about it: we can, by relating it to other concepts like belief, desire, cause and action. Nor does the indefinability of truth imply that the concept is mysterious, ambiguous, or untrustworthy [Davidson, 1996, 21].

Williams, 2002, 63

Two comments about this. First, the genealogist is not interested in *definition* in any case. How is it relevant to the possibility of a genealogy of truth that a different project, one that the genealogist has in any case rejected, is impossible? Second, while truth may indeed be part of a package, it would be absurd to suggest that the package as whole does not have a history. Our ancestors didn't always play the game of giving and asking for reasons, of course.

In the paragraph preceding the one from which Williams quotes, Davidson puts his point more generally.

For the most part, the concepts philosophers single out for attention, like truth, knowledge, belief, action, cause, the good and the right, are the most elementary concepts we have, concepts without which (I am inclined to say) we would have no concepts at all. Why then should we expect to be able to reduce these concepts definitionally to other concepts that are simpler, clearer, and more basic? We should accept the fact that what makes these concepts so important must also foreclose on the possibility of finding a foundation for them that reaches deeper into bedrock. Davidson, 1996, 20–21

In the case of knowledge, Williams does not seem to regard Davidson's prohibition on reductive definition as an obstacle to Craig's genealogy. Why should truth be different?

Turning to Davidson himself, it is easy to read his presentations of the triangulation argument as proto-State of Nature stories, in Craig's sense. Consider this version from 'Seeing Through Language', for example.

The sharing of responses to stimuli found similar allows an interpersonal element to emerge: creatures that share responses can correlate each other's responses with what they are responses to. Person A responds to person B's responses to situations both A and B find similar. A triangle is thus set up, the three corners being A, B, and the objects, events, or situations to which they mutually respond. This elaborate, but commonplace, triangular interaction between creatures and a shared environment does not require thought or language; it occurs with great frequency among animals that neither think nor talk. Birds and fish do it as well as monkeys, elephants, and whales.

Davidson, 1997, 140

So far, this sounds very much like an account of the State of Nature of our pre-linguistic ancestors. Davidson continues like this:

What more is there to linguistic communication and developed thought? The answer is, I think, two things that depend on the basic triangle, and emerge from it. The first is the concept of error, that is, appreciation of the distinction between belief and truth. The interactions of the triangle do not in themselves automatically generate this appreciation, as we see from the example of simple animals, but the triangle does make room for the concept of error (and hence of truth) in situations in which the correlation of reactions that have been repeatedly shared can be seen by the sharers to break down; one creature reacts in a way previously associated by both creatures with a certain sort of situation, but the other does not. This may simply alert the non-reactor to an unnoticed danger or opportunity, but if the anticipated danger or opportunity fails to materialize, a place exists for the notion of a mistake. We, looking on, will judge that the first creature erred. The creatures themselves are also in a position to come to the same conclusion. If they do, they have grasped the concept of objective truth.)

With the second, final, step, we move in a circle, for we grasp the concept of truth only when we can communicate

the contents—the propositional contents—of the shared experience, and this requires language. The primitive triangle, constituted by two (and typically more than two) creatures reacting in concert to features of the world and to each other's reactions, thus provides the framework in which thought and language can evolve. Neither thought nor language, according to this account, can come first, for each requires the other. This presents no puzzle about priorities: the abilities to speak, perceive, and think develop together, gradually. We perceive the world through language, that is, through having language.

Davidson, 1997, 140–141

Read from Craig's genealogical perspective, this suggests some questions. Why do we have these concepts of truth and error? What new job do they do, in a State of Nature of this kind?

Davidson often emphasises that the concepts of truth and error don't emerge automatically from the triangular arrangement.

Here is my thesis: an interconnected triangle such as this (two lionesses, one gazelle) constitutes a necessary condition for the existence of conceptualization, thought, and language. It makes possible objective belief and the other propositional attitudes. [...] The simple model I have described does not, to be sure, suggest where the norm lies when the creatures differ in their responses [...] The point is not to identify the norm, but to make sense of there being a norm, and this has been done if we can point to the difference between the preponderance of cases where the creatures respond alike and the deviant cases where they diverge.

Davidson, 2001d, 9

Very well, but instead of asking 'where the norm lies'—whatever Davidson means by that—let's ask the genealogist's questions. What difference does it make to *invent* a norm of this kind? And what did our ancestors need to *do*, to 'institute'¹⁴ and sustain it?

Later in the same paper Davidson continues like this:

We must remember, of course, that if going on in the right way means only going on as the teacher does, this is not enough in itself to constitute mastery of word or concept; there remains

¹⁴ To borrow a term from Brandom.

the grasp of truth conditions, the awareness of the possibility of error. Nor will punishing or correcting what the teacher sees as error suffice, for the same reason. We can say that in the potential gap between how the learner goes on and how the teacher goes on there is room to introduce the idea of error, not only in the mind of the teacher, for we are assuming it is *there*, but also in the mind of the learner. But how that room gets filled is probably beyond our powers to tell. Our failure to provide an analysis of the concept of error or, as we could as well say, of the concept of objectivity [...] does not mean there is something hopelessly mysterious about these concepts; it only reflects the fact that intentional phenomena cannot be reduced to something simpler or different. Nor does our failure mean that true grasp of concepts as opposed to mere following of routines does not reveal itself in behavior. What we cannot do is say, in non-intentional terms, in what the difference consists. Davidson, 2001d, 17

Again, my suggestion is that Davidson is looking in the wrong place. Let's grant that there is no informative *reduction* of the norms of truth and error to something 'simpler or different'. That doesn't prevent us from asking what proto-speakers need to *do*, in order to establish such a norm in their practice. I propose that what they need to do, most basically, is adopt the practice of *commending* or *criticising* fellow speakers with whom they agree or disagree, respectively. In other words, most basically, they need to treat disagreement as a sign of *fault*, on another speaker's part.

More needs to be said, of course. What is the proto-linguistic practice within which such a practice of blame and praise could develop?¹⁵ And what does disagreement mean, if it not to be explained, circularly, as recognition that another speaker has said something *false*, by one's own lights. These are good questions, but Davidson's own elaboration of triangulation provides some of the answers, as we have seen.¹⁶ Note that praise and blame themselves do not need to be invented, at this point in the story. As communal creatures, we surely had social carrots and sticks already. The task was simply to put them to use in a new way.

Why should Davidson object to such an approach? This is what he says about his own methodology, in 'The Folly of Trying to Define Truth'.

¹⁵ Here the work of Dorit Bar-On (2013, 2016), on the continuities between animal communication and human language, would be a good place to start.

¹⁶ I discussed some of these issues about disagreement in (Price, 1990).

I want to describe what I take to be a fairly radical alternative to the theories I have been discussing and (with unseemly haste) dismissing. What I stress here is the *methodology* I think is required rather than the more detailed account I have given elsewhere. The methodology can be characterized on the negative side by saying it offers no definition of the concept of truth, nor any quasi-definitional clause, axiom schema, or other brief substitute for a definition. The positive proposal is to attempt to trace the connections between the concept of truth and the human attitudes and acts that give it body.

Davidson, 1996, 35

The approach I have suggested certainly agrees with Davidson 'on the negative side'. It agrees on the positive side, too, in a sense. It is certainly proposing 'connections between the concept of truth and the human attitudes and acts that give it body'. Yet it is doing this in a vocabulary that Davidson does not appear to have in his repertoire.

A cheeky suggestion would be that Davidson was a product of the wrong Cambridge, at the wrong time, somehow falling between the pragmatist pillars in each. He was no stranger to Ramsey, of course. He coined the term 'Ramsey effect' (1999, 32), for the discovery that Ramsey already solved the problem to which one has just devoted so much thought. Yet he seems to have missed Ramsey's brand of pragmatic genealogy, which throws such useful though non-reductive light on concepts such as truth, probability and causation—in other words, on some of the very concepts that Davidson lists as 'so important [that we] must [...] foreclose on the possibility of finding a foundation for them that reaches deeper into bedrock' (1996, 20–21).

This question about Davidson's apparent blindness to genealogy deserves a more thorough treatment than I can give it here. Writers who know Davidson better than me might be able to reach deeper into his own bedrock, to throw some light on the matter. But for the moment, concerning truth, I conclude that Davidson, like Rorty and Williams himself, could have benefitted from a deeper immersion in Ramsey's Cambridge pragmatism.

References

- Bar-On, Dorit (2013). "Expressive Communication and Continuity Skepticism." *Journal of Philosophy* 110, 293-330.
- Bar-On, Dorit (2016). "Sociality, Expression, and This Thing Called Language." *Inquiry* 59, 56-79.
- Brandom, Robert B. (ed.) (2000). *Rorty and His Critics*. Malden: Blackwell.
- Craig, Edward (1990). *Knowledge and the State of Nature*. Oxford: Clarendon Press.
- Davidson, Donald (1974). "On the Very Idea of a Conceptual Scheme." *Proceedings and Addresses of the American Philosophical Association*, 47: 5-20. Reprinted in Davidson 2001b, 183-198. Page references are to the latter version.
- Davidson, Donald (1990). "The Structure and Content of Truth." *Journal of Philosophy* 87, 279-328.
- Davidson, Donald (1991). "Three Varieties of Knowledge." In: A. Phillips Griffiths (ed.), *A.J. Ayer Memorial Essays: Royal Institute of Philosophy Supplement*, 30. Cambridge: Cambridge University Press. Reprinted in Davidson 2001c, 205-220. Page references are to the latter version.
- Davidson, Donald (1996). "The Folly of Trying to Define Truth." *Journal of Philosophy* 93 (6):263-278. Reprinted in Davidson (2005), 19-38. Page references are to the latter version. <https://doi.org/10.1093/019823757X.003.0002>
- Davidson, Donald (1997). "Seeing Through Language." [https://philpapers.org/a/search.pl?pub=991Royal Institute of Philosophy Supplement 42:15-27 \(1997\)](https://philpapers.org/a/search.pl?pub=991Royal+Institute+of+Philosophy+Supplement+42:15-27+(1997).). Reprinted in Davidson (2005), 127-142. Page references are to the latter version. <https://doi.org/10.1093/019823757X.003.0009>
- Davidson, Donald (1999). "Intellectual Autobiography of Donald Davidson." In: L. Hahn (ed.), *The Philosophy of Donald Davidson*. Chicago & LaSalle, IL: Open Court, 3-70.
- Davidson, Donald (2001a). *Essays on Actions and Events*. Oxford: Clarendon Press, 2nd edn.
- Davidson, Donald (2001b). *Inquiries into Truth and Interpretation*. Oxford: Clarendon Press, 2nd edn.
- Davidson, Donald (2001c). *Subjective, Intersubjective, Objective*. Oxford: Clarendon Press.
- Davidson, Donald (2001d). "Externalisms." In: Kotatko, Pagin, and Segal G. (2001), 3-16.
- Davidson, Donald (2005) *Truth, Language, and History: Philosophical Essays Volume 5* Oxford: Clarendon Press.
- Dummett, Michael (1959). "Truth." *Proceedings of the Aristotelian Society* 59, 141-62. Reprinted in *Truth and Other Enigmas*. London: Duckworth, 1978, pp. 1-24. Page references are to the latter version.
- Hahn, Lewis Edwin (ed.) (1999). *The Philosophy of Donald Davidson*. Open Court, Chicago.

- Gert, Joshua, (ed.) (2023). *Neopragmatism: Interventions in First-order Philosophy*. New York: Oxford University Press.
- Gibbard, Allan (2015). "Global Expressivism and the Truth in Representation." In: Gross S, Tebben N., Williams M. (eds.), *Meaning without representation: essays on truth, expression, normativity, and naturalism*. New York, Oxford University Press, 210–223.
- Kotatko, P., Pagin, P. and Segal G. (eds.) (2001). *Interpreting Davidson*. Stanford: CSLI Publications.
- Jackson, Frank, and Philip Pettit (1998). "A Problem for Expressivism." *Analysis*, 58, 239–251.
- Millikan, Ruth (1989). "Biosemantics." *Journal of Philosophy* 86, 281–97.
- Millikan, Ruth (2009). "Biosemantics." In: McLaughlin B., Beckermann A., Walter S. (eds.), *Oxford Handbook of Philosophy of Mind*, Oxford, Oxford University Press, 394–406.
- Price, Huw (1983). "Could a Question be True?: Assent and the Basis of Meaning." *The Philosophical Quarterly* 33, 354–64.
- Price, Huw (1988). *Facts and the Function of Truth*. Oxford: Blackwell.
- Price, Huw (1990). "Why 'Not'?" *Mind* 99, 221–238.
- Price, Huw (1991). "Two Paths to Pragmatism." In: Peter Menzies (ed.), *Response-Dependent Concepts*. Canberra: Philosophy Program, RSSH, ANU, 46–82. Reprinted in Price (2011), 80–111.
- Price, Huw (1997). "Naturalism and the Fate of the M-Worlds." *Proceedings of the Aristotelian Society, Supp. Vol.*, LXXI, 247–267. Reprinted in Price (2011), 132–147.
- Price, Huw (1998a). "Three Norms of Assertibility, or How theMOA Became Extinct." *Philosophical Perspectives* 12, 41–54.
- Price, Huw (1998b). "Ramsey, Reference and Reductionism." Retrieved August 13, 2022 from: <https://philpapers.org/rec/PRIRRA-2>
- Price, Huw (1999). Location, Location, Location. Presidential Address to the Australasian Association of Philosophy, Melbourne. Retrieved November 24, 2022 from: <https://philpapers.org/go.pl?aid=PRILLL>.
- Price, Huw (2003). "Truth as Convenient Friction." *Journal of Philosophy* 100, 167–90. Reprinted in Price (2011), 163–183. Page references are to the latter version.
- Price, Huw (2004). "Naturalism Without Representationalism." In: D. Macarthur and M. de Caro (eds.), *Naturalism in Question*. Cambridge, MA: Harvard University Press, 71–88. Reprinted in Price (2011), 184–199.
- Price, Huw (2007). "Causal Perspectivalism." In: Price and Corry (2007), 250–292.
- Price, Huw (2011). *Naturalism without Mirrors*. New York: Oxford University Press.
- Price, Huw (2017). "Epilogue: Ramsey's Ubiquitous Pragmatism." In: Cheryl Misak and Huw Price (eds.), *The Practical Turn: Pragmatism in the British Long Twentieth Century*. Oxford, Oxford University Press, 149–162.

- Price, Huw (2023a). "Gibbard on Quasi-Realism and Global Expressivism". *Topoi* 42, 683–697.
- Price, Huw (2023b). "Time for Pragmatism." In: Gert (2023b), 23–69.
- Price, Huw and Corry, Richard (eds.) (2007). *Causation, Physics, and the Constitution of Reality: Russell's Republic Revisited*. Oxford, Oxford University Press.
- Ramberg, Bjørn Torgrim (2000). "Post-Ontological Philosophy of Mind: Rorty versus Davidson." In: Brandom, Robert (ed.), *Rorty and his Critics*. Oxford: Blackwell, 351–370.
- Ramberg, Bjørn Torgrim (2001). "What Davidson Said to the Sceptic, or Anti-representationalism, Triangulation, and the Naturalization of the Subjective." In: Kotatko, Pagin, and Segal G. (2001), 213–236.
- Ramberg, Bjørn Torgrim (2022). "Rorty as Mentor." In: Müller, M. (ed.) *Handbuch Richard Rorty*. Wiesbaden, Springer vs, 115–117.
- Rorty, Richard (1979). *Philosophy and the Mirror of Nature*. Princeton, Princeton University Press.
- Rorty, Richard (1995). "Is Truth a Goal of Enquiry? Donald Davidson versus Crispin Wright." *Philosophical Quarterly*, 45, 281–300. Reprinted in *Truth and Progress: Philosophical Papers Vol. 3*. Cambridge University Press, 1998, 19–42. Page references are to the latter version.
- Rorty, Richard (1999). "Davidson's Mental-Physical Distinction." In: Hahn (ed.) (1999), 575–594.
- Rorty, Richard and Price, Huw (2010). "Exchange on 'Truth as Convenient Friction'." In: Mario De Caro and David Macarthur (eds.), *Naturalism and Normativity*. New York: Columbia University Press, 253–262.
- Williams, Bernard (2002). *Truth and Truthfulness: An Essay in Genealogy*. Princeton: Princeton University Press.
- Wright, Crispin (1992). *Truth and Objectivity*. Cambridge, Mass: Harvard University Press.
- Wright, Crispin (2021). "Alethic Pluralism, Deflationism, and Faultless Disagreement." *Metaphilosophy* 52: pp. 432–48.