Truth as Convenient Friction†

Huw Price

[huw@mail.usyd.edu.au]

1. Introduction

In a recent paper, Richard Rorty begins by telling us why pragmatists such as himself are inclined to identify truth with justification:

> "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."[1]

Rorty goes on to discuss the claim, defended by Crispin Wright, that truth is a normative constraint on assertion. He argues that this claim runs foul of this principle of no difference without a practical difference:

> "The need to justify our beliefs to ourselves and our fellow agents subjects us to norms, and obedience to these norms produces a behavioural pattern that we must detect in others before confidently attributing beliefs to them. But there seems to be no occasion to look for obedience to an additional norm – the commandment to seek the truth. For – to return to the pragmatist doubt with which I began – obedience to that commandment will produce no behaviour not produced by the need to offer justification." (p. 26)

Again, then, Rorty appeals to the claim that a commitment to a norm of truth rather than a norm of justification makes no behavioural difference. In principle, this is an empirical claim, decidable by comparing the behaviour of a community of

---

† This is a draft paper, as presented under the title "No Conversation Without Representation" at a conference in honour of Richard Rorty, at the Humanities Research Centre, ANU, 9–11 July, 1999. Please don't cite the paper without permission.

[1] "Is truth a goal of enquiry? Donald Davidson versus Crispin Wright", Philosophical Quarterly 45(1995), 281-200; reprinted in Rorty's Truth and Progress: Philosophical Papers Vol. 3, CUP, 1998, 19-42; at p. 19.

thoroughgoing realists to that of a community of thoroughgoing pragmatists. It could be done by experiment, by creating a community of realists to compare with pragmatist controls (or, as in Australia, the other way around).

I think that the experiment would show that claim is unjustified, indeed false: there is an important and widespread behavioural pattern which results from the fact that speakers take themselves to be subject to such an additional norm. In fact, it is behavioural pattern so central to what we regard as a worthwhile human life that no reasonable person—indeed, few research ethics committees—would knowingly condone the experiment. Moreover, it is a behavioural pattern which Rorty of all people cannot afford to dismiss as a pathological and dispensable by-product of bad philosophy: in a word, it is conversation itself.

In other words, I want to show that in order to account for ordinary conversational practice, we must allow that speakers take themselves to be governed by a norm stronger than that of justification—a norm which speakers acknowledge they may fail to meet, even if their claims are well-justified. It turns out to be helpful to distinguish three norms, in order of increasing strength: roughly, sincerity, justification and truth itself. In order to throw into relief the crucial role of the third norm in conversation, I shall contrast conversation with some non-conversational (and non-assertoric) uses of language. In these cases, I'll argue, the two weaker norms still apply. Moreover, it turns out that some of the basic functions of assertoric discourse could be fulfilled in an analogous way, by a practice which lacked the third norm. But it will be clear, I hope, that that practice would not be conversation as we know it. What is missing – what the third norm provides – is the automatic and quite unconscious sense of engagement in common purpose that distinguishes conversation 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.[2]

---

[2]If private cognition depends on public cognition, then truth plays the same role, at second hand, in the private sphere. I think this is a plausible view, but I won't try to defend it here.

To switch to a Rylean metaphor, my view is thus that truth supplies conversation with its essential esprit de corps. As the metaphor is meant to suggest, what matters is that would-be conversationalists think that there is such a norm – that they take themselves to be governed by it – not that their view be confirmed by science or metaphysics. (Science has already done its work, in pointing out the function of the thought in the lives of creatures like us.) This may suggest that a commitment to truth is like a commitment to theism, an analogy which Rorty himself draws, against Crispin Wright, in the paper with which I began. However, there are several important differences.

First, there are no devastating behavioural consequences in giving up theism, at least in this life. But if I am right, then the consequences of giving up truth would be rather more serious. It would silence the conversation of mankind – though "silence" may be the wrong word, if the alternative is disengaged chatter. (Even more seriously, I suspect it would silence our own "internal" conversations, but I won't discuss this consequence here.)

Second, it is doubtful whether giving up truth is really an option open to us. I suspect that people who think it is an option haven't realised how deeply embedded the idea of truth is in linguistic practice, and therefore underestimate the extent of the required change.[3]

Third, and perhaps most interestingly of all, the issue of the status of truth is enmeshed with the terms of the problem, in a way which is quite uncharacteristic of the theism case. Metaphysical conclusions tend to be cast in semantic vocabulary: theism is in error in virtue of the fact that its claim are not true; that its terms fail to refer. For this reason, it is uniquely difficult to formulate a meaningful antirealism about the semantic terms themselves. In my view the right response to this is not to think (with Boghossian[4]) that we thereby have a transcendental argument for semantic realism. Without an intelligible denial, realism is no more intelligible than antirealism. The right response – as

---

[3]In a recent review article, Jonathon Rée accuses Rorty of underestimating the practical inflexibility of our admittedly contingent practices.

[4]"The Status of Content", Philosophical Review, 99(1990), 157–184.

Rorty himself in any case urges – is to be suspicious about the realist/antirealist debate itself. (A realist could object that a commitment to the third norm might be useful and yet in error, but Rorty can't. It is fair for him to object against Wright that this commitment might be like theism, because Wright takes metaphysics seriously. By Wright's professed standards, then, the theism objection poses a real threat.)

The interesting thing is that Rorty ties rejection of the realist/antirealist debate to rejection of a notion of truth distinct from justification, and of the idea of representation. I think this is the wrong path to the right conclusion. We should reject the metaphysical stance not by rejecting truth and representation, but by recognising that in virtue of the most plausible story about the function and origins of these notions, they simply don't sustain that sort of metaphysical weight.

Rorty describes himself as oscillating between pragmatism about truth—i.e., "trying to reduce truth to justification" (p. 21)—and deflationism. My view is a defence of the latter strategy, though not by the familiar route. My account of truth is deflationary in that it does not regard truth as a "substantial" property, about the nature of which there is an interesting philosophical issue. Like more familiar deflationists such as Quine and Horwich, I think that the right approach to truth is to investigate its function in human discourse: to ask what difference it makes to us to have such a concept. Unlike such deflationists, however, I don't think the right answer to this question is that truth is merely a grammatical device for disquotation. It has a more important function, which requires that it be treated as a norm. But like other deflationists, I think that there is no further question of interest to philosophy, once the question about function has been answered.

In one sense, my view of truth is itself pragmatist, for it explicates truth in terms of its role in practice. (This is also true of Horwich's deflationism.) In another sense, it conflicts with pragmatism, for it opposes the proposal that we identify truth with justification. In my view, this contrast reflects a deep tension within pragmatism. From Peirce and James on, pragmatists have often been unable to resist the urge to join their opponents in asking "What is truth?" (Indeed, the pragmatist position as a whole is often

characterised in terms of its answer to this question.) Pragmatism thus turns its back on alternative paths to philosophical illumination about truth, even though these alternative paths—explanatory and genealogical approaches—are at least compatible with, if not mandated by, the pragmatist doctrine that we understand problematic notions in terms of their practical significance.

Rorty himself is well aware of this tension in pragmatism. In "Pragmatism, Davidson and Truth", for example, he notes that James is less prone than Peirce to try to answer the "ontological" or reductive question about truth, and suggests that Davidson may be thought of as a pragmatist in the preferable non-reductive sense. (Robert Brandom makes a similar point in "Pragmatism, Phenomenalism, and Truth Talk", Midwest Studies, XII.) As he oscillates between pragmatism and deflationism, then, Rorty himself is at worst only intermittently subject to this craving for an analysis of truth. All the same, it seems to me that he is never properly aware of the range of possibilities for non-reductive pragmatism about truth. In particular, he is not properly aware of the possibility that such a pragmatism might find itself explaining the fact that the notion of truth in ordinary use is (and ought to be) one that conflicts with the identification of truth with justification: a normative goal of inquiry, stronger than any norm of justification, of the very kind that realists about truth – opponents both of pragmatism and of disquotationalism – mistakenly sought to analyse. In other words, Rorty seems to miss the possibility that the right thing for the explanatory pragmatist to say might be that truth is a goal of inquiry distinct from norms of justification, and that the realist's mistake is to try to analyse this normative notion, rather than simply to map its function and possible genealogy. It is that possibility that I want to defend.

2. Falsity and lesser evils

As I have said, I want to argue that truth plays a crucial role as a norm of assertoric discourse. It is not the only such norm, however, and a good way to highlight the distinctive role of truth is to distinguish certain weaker norms, and to imagine a linguistic

practice which had those norms but not truth. By seeing what such a practice lacks, we see what truth adds.

There are at least two weaker norms of assertion, in addition to any distinctive norm of truth. The weakest relevant norm seems to be that embodied in the principle that it is prima facie appropriate to assert that p when and only when one believes that p—prima facie, because of course many other factors may come into play, in determining the appropriateness of a particular assertion in a particular context. Let's call this the norm of subjective assertibility. (This corresponds to a common use of the term "assertibility condition", as for example when it is said that the subjective assertibility condition for the indicative conditional "If p then q" is a high conditional credence in q given p.)

The norm is perhaps best characterised in negative form:

(Subjective assertibility)

One is incorrect to assert that p if one does not believe that p.

The easiest way to see that this norm has very little to do with truth is to note that it is exactly analogous to norms which operate with respect to utterances which we don't take to be truth-apt. It is prima facie appropriate to request a cup of coffee when and only when one wants a cup of coffee, but this doesn't show that requests or desires are subject to a norm of truth. In effect, this norm is simply that of sincerity, and some such norm seems to govern much conventional behaviour. Conventions often depend on the fact that communities censure those who break them.

The second norm is that of (personal) warranted assertibility. Roughly, "p" is warrantedly assertible by a speaker who not only believes that p, but is justified in doing so. The qualification "personal" recognises the fact that there are different kinds and degrees of warrant or justification, some of them more subjective than others. For example, is justification to be assessed with reference to a speaker's actual evidence as she (presently?) sees it, or by some less subjective lights? For the moment, for a degree of definiteness, let us think of it in terms of subjective coherence—a belief is justified if

supported by a speaker's other current beliefs. This is what I shall mean by personal warranted assertibility.

Again, this second norm is usefully characterised in negative form:

(Personal warranted assertibility)

> One is incorrect to assert that p if one does not have adequate (personal) grounds for believing that p.

A person who meets both the norms just identified may be said to have done as much as possible, by her own current lights, to ensure that her assertion that p is in order. Obviously, realists will say that her assertion may nevertheless be incorrect: subjective assertibility and (personal) warranted assertibility do not guarantee truth. To an extent, moreover, most pragmatists are likely to agree. Few people who advocate reducing truth to (or replacing truth by) a notion of warranted assertibility have personal warranted assertibility in mind. Rather, they imagine some more objective, community-based variant, according to which a belief is justified if it coheres appropriately with the other beliefs of one's community. If we call this communal warranted assertibility, then the point is that we can make sense of a gap between the personal and communal notions. A belief may be justified in one sense but not the other.

Pragmatists and realists may thus agree that there is a normative dimension distinct from subjective assertibility and personal warranted assertibility—an assertion may be wrong, despite meeting these norms. The fact that philosophers agree on this point does not yet show that the norm in question need be marked in ordinary discourse. In principle, it might properly remain a theoretical notion, useful in second-order reflection on linguistic practice but unnecessary in ordinary talk about other matters. In practice, however, there is a very good reason why it should not remain theoretical. Unless individual speakers recognise such a norm, the idea that they might improve their views by consultation with the wider community is simply incoherent to them. (It would be as if we gave a student full marks in an exam, and then told him that he would have done better if his answers had agreed with those of other students.)

It may seem that as yet, this argument doesn't favour realism over pragmatism. If the normative standard an individual speaker needs to acknowledge is that of the community as a whole, there is as yet no pressure to a notion of truth beyond community-wide warranted assertibility. But what constitutes the relevant community? At any given stage, isn't the relation of a given community to its possible present and future extensions just like that of the individual to her community? If so, then the same argument applies at this level. At each stage, the actual community needs to recognise that it may be wrong by the standards of some broader community. (Cf. Rorty, p. 22: "For any audience, one can imagine a better-informed audience.")

The pragmatist might now seem obliged to follow Peirce, in identifying truth with warranted assertibility in the ideal limit of inquiry. The useful thing about this limit, in this context, is that it transcends any actual community. But in my view, as I'll explain below (and as Rorty in some moods already case agrees), a better move for a pragmatist is to resist the pressure to identify truth with anything—in other words, simply to reject the assumption that an adequate philosophical account of truth needs to answer the question "What is truth?" Better questions for a pragmatist to ask are the explanatory ones: Why do we have such a notion? What job does it do in language? What features does it need to have to play this role? And how would things be different if we didn't have it?

For the moment, we have the beginnings of an answer to the last question. If we didn't have a normative notion in addition the norms of subjective assertibility and personal warranted assertibility, the idea that we might improve our commitments by seeking to align them with those of our community would be simply incoherent. I'll call this the passive argument for the third norm—'passive', because it doesn't yet provide an active or causal role for a commitment to truth. Later, I'll argue that the third norm not only creates the conceptual space for argument, in this passive sense, but actively encourages speakers to participate.

3. The third norm in focus

The best way to bring the third norm into focus is again to consider its negative form:

(Truth)

If not-p then it is incorrect to assert that p.

The crucial point is that there is a norm of assertion which a speaker may fail to meet, even if she does meet the norms of subjective assertibility and (personal) warranted assertibility. We judge a speaker wrong, incorrect, mistaken, when we are prepared to make a contrary assertion, even if we are in no doubt that she is sincere, and in possession of the kind of evidence that would lead any reasonable person to make the same mistake.

One of the reasons why this third norm is hard to distinguish from the two weaker norms of assertibility is that when we apply it in judging a fellow speaker right or wrong, the basis for our judgement lies in our own beliefs and evidence. It is not as if we are in a position to make the judgement from the stance of reality itself, as it were. I think this can make it seem as if application of this norm involves nothing more than re-assertion of the original claim (in the case in which we judge it correct), or assertion of the negation of the original claim (in the case in which we judge it incorrect). Construed in these terms, our response contains nothing problematic for orthodox disquotational versions of the deflationary view, of course. Re-assertion of this sort is precisely one of the linguistic activities which disquotational truth facilitates. Construed in these terms, then, there is no need for truth to be a distinct norm.

But our response is not merely re-assertion, or assertion of the negation of the original claim. If it were, it would involve no commendation or criticism of the original utterance. This is hard to see, but the crucial point is that we can imagine a linguistic practice which allowed re-assertion and contrary assertion, but without this third normative dimension. That is, we can imagine a linguistic community who use sentences to express their beliefs, and have a purely disquotational truth predicate, but for whom disagreements have no normative significance, except in so far as it is related to the weaker norms of assertibility.

What we need is the idea of a community who take an assertion—or rather the closest thing they have to what we call an assertion—to be merely an expression of the speaker's opinion. The relevant idea is familiar in the case of expressions of desires and preferences. It is easy to imagine a community—we are at least close to it ourselves—who have a language in which they give voice to psychological states of these kinds, not by reporting that they hold them (which depends on assertion), but directly, in conventional linguistic forms tailored specifically for this purpose.

Think of a community who use language only for expressing preferences in restaurants, for example. (Perhaps the development of such a restricted language from scratch is incoherent, but surely we might approach it from the other direction: imagine a community of dedicated lunchers, whose language atrophies to the bare essentials.) In this community we would expect a norm analogous to subjective assertibility: essentially, a normative requirement that speakers use these conventional expressions sincerely. Less obviously, such a practice might also involve a norm analogous to personal warranted assertibility. In other words, expressed preferences might be criticised on the grounds that they were not well-founded, by the speaker's own lights (for example, on the grounds they did not cohere with the speaker's other preferences and desires). However, in this practice there need be no place for a norm analogous to truth—no idea of an objective standard, over and above personal warranted assertibility, which preferences properly aim to meet.

At least to a first approximation, we can imagine a community who treat expressions of beliefs in the same way. They express their beliefs—i.e., the kind of behavioural dispositions which we would characterise as beliefs—by means of a speech act we might call the merely-opinionated assertion ("MOA", for short). These speakers—"Mo'ans", as I called them in another paper[5]—criticise each other for insincerity and for lack of coherence, or personal warranted assertibility. But they go no further than this. In particular, they do not treat a disagreement between two speakers as

---

[5] 'Three norms of assertibility, or how the MOA became extinct', in J. Tomberlin (ed.) Philosophical Perspectives, Vol. XII, Language, Mind, and Ontology, Blackwell, 1998, 241–54. Some sections of the present paper draw heavily on this paper.

an indication that, necessarily, one speaker or other is mistaken—in violation of some norm. On the contrary, they allow that in such a case it may turn out that both speakers have spoken correctly, by the only two standards the community takes to be operable. Both may be sincere, and both, in their own terms, may have good grounds for their assertion.

This speech community could quite well make use of deflationary truth, for example as a device to facilitate agreement with an expression of opinion made by another speaker. "That's true" would function much like "Same again", used in a bar or restaurant. Just as "Same again" serves to indicate that one has the same preference as a previous speaker, "That's true" would serve to indicate that one holds the same opinion as the previous speaker. The crucial point is that if the only norms in play are subjective assertibility and personal warranted assertibility, introducing disquotational truth leaves everything as it is. It doesn't import a third norm.

The difficulty we have in holding on to the idea of such a community stems from our almost irresistible urge to see the situation in terms of our own normative standards. There really is a third norm, we are inclined to think, even if these simple creatures don't know it. When they make incompatible assertions, at least one of them must be objectively incorrect, even if by their lights they both meet the only norms they themselves recognise. (I think even pragmatists will be inclined to say this, even though they want to equate the relevant kind of incorrectness not with falsity but with lack of some kind of justification more objective than that of personal warrant.) But the point of the story is precisely to bring this third norm into sharp relief, and hence I am quite happy to allow challenges to the story on these grounds, which rely on the very conclusion I want to draw: For us, there is a third norm. But why is that so? Where does the third norm come from? What job does it do—what difference does it make to our lives? What features must it have in order to do this job?

4. What difference does the third norm make?

Let's return to the Mo'ans, and their merely-opinionated assertions. Recall that Mo'ans use linguistic utterances to express their beliefs (as well as other psychological states, such as preferences and desires). Where they differ from us is in the fact that they do not take a disagreement between two speakers in this belief-expressing linguistic dimension to indicate that one or other speaker must be at fault. They recognise the possibility of fault consisting in failure to observe one of the two norms of subjective assertibility or personal warranted assertibility, but lack the idea of the third norm, that of truth itself. This shows up in the fact that by default, disagreements tend to be of a no-fault kind (in the way that expression of different preferences often are for us).

Think of the Mo'ans as speakers of a proto-language. How are we to understand what happens when these speakers hit upon the third norm of assertibility—the notion of truth—and the MOA becomes extinct? I think it is important to realise that there are two quite different possibilities. One possibility is that Mo'ans gradually come to realise that there is an important pre-existing property that the psychological attitudes they use MOAs to express may have or lack: perhaps the property of corresponding to how things are in the world, or—as we would put it—of being <u>true</u>; or perhaps the property of being what their opinions are fated to converge on in the long run. Perhaps this property is in some sense essentially normative. If not, then it might at least be such that the sensible Mo'ans come to recognise its importance, and treat it as a norm—in other words, they come to adopt the convention that an expression of a belief is at fault, in so far as the belief fails to possess this property. Call this the <u>substantial</u> account of how the MOA becomes extinct.

The second alternative is quite different. Suppose there is no substantial, objective, property of this kind, which the Mo'ans' belief-like behavioural dispositions either have or lack. Nevertheless, it might turn out to be very much to the Mo'ans' advantage to behave as if there were such a property. As it turns out, it isn't difficult to adopt this pretence. The practice Mo'ans need to adopt is exactly the same as that required by the previous alternative. They simply need to ensure that when they believe that p, they be

prepared not only to assert (in the old MOA sense) that p, but also to ascribe fault to anyone who asserts not-p, independently of any grounds for thinking that that person fails one of the first two norms of assertibility. In other words, the usage rule for this imaginary norm is exactly what the corresponding rule would be according to the first story—which is hardly surprising, for it is effectively the disquotational schema, transformed into the rule that one should be prepared to assert that p is correct, if and only if one is prepared to assert that p.

Why might the invention of such a norm be useful? We already have one answer, in the passive argument: without a norm stronger than that of warranted assertibility for me, or for us, the idea of improving my, or our, commitments is incoherent. Since it is impractical and undesirable to put bounds in advance on where improvement might end, we need a norm stronger than that of warranted assertibility for any actual community. (Of course, this doesn't yet show that we need something more than Peircean ideal assertibility, but one thing at a time.)

However, we can do better than the passive argument. The third norm doesn't just hold open the conceptual space for the idea of improvement; it positively encourages such improvement, by motivating speakers who disagree to try to resolve their disagreement. Without the third norm, differences of opinion would simply slide past one another. Differences of opinion would seem as inconsequential as differences of preference. With the third norm, however, disagreement automatically becomes normatively loaded. The third norm makes what would otherwise be no-fault disagreements into unstable social situations, whose instability is only resolved by argument and consequent agreement—and it provides an immediate incentive for argument, in that it holds out to the successful arguer the reward consisting in her community's positive evaluation of her dialectical position. If reasoned argument is generally beneficial—beneficial in a long-run sense—then a community of Mo'ans who adopt this practice will tend to prosper, compared to a community who do not.

I'll call this the active argument for the third norm. In effect, it contends that the fact that speakers take their belief-expressing utterances to be subject to the third norm

plays a causal, carrot-and-stick role in encouraging them to settle their differences, in cases in which initially they disagree.

Much more needs to be said about possible advantages of such a mechanism for resolving disagreements—about its long-run advantages, for example, both compared to the case in which there is no such mechanism and compared to the case in which there is some different mechanism (such as deference to social rank). For immediate purposes, however, my claim does not depend on all this further work. For the present, my claim is simply that truth does play the role of this third norm, in providing the friction essential to "factual" conversation as we know it. (I also claim, roughly, that this is the most interesting fact about truth, from a philosophical perspective.) In principle, this claim could be true, even though the practice in question were not advantageous. In principle, truth, and with it factual conversation, could turn out to be a bad thing for the species, biologically considered. No matter. It would still be true that we wouldn't have understood truth until we understood its role in this debilitating practice.

5. Peirce regained?

Now to the question deferred above. Why does the third norm need to be anything other than a more-than-merely-personal notion of justification? In particular, couldn't it be a Peircean flavour of ideal warranted assertibility? I have several responses to this suggestion.

First, I think that the suggestion is mistakenly motivated. As I said in the introduction, I think it stems from the tendency, still too strong in Peirce, to ask the wrong question about truth. If we think that the philosophical issue is "What is truth?", then naturally we'll want to find an answer—something with which we may identify truth. Then, given standard objections to metaphysical answers, it is understandable that Peirce's alternative should seem attractive. But the attraction is that of methadone compared to heroin—far better, surely, from a pragmatist's point of view, to rid ourselves of the craving for analysis altogether. To do this, we need to see that the basic philosophical needs that analysis seemed to serve can be met in another mode altogether:

by explanation of the practices, rather than reduction of their objects. (Moreover, the explanatory project has the potential to allow us realist truth without the metaphysical disadvantages. The apparent disadvantages of realist truth emerge in the light of the reductive project; for it is from this perspective that it seems mysterious what truth could be. If we no longer feel obliged to ask the question, we won't be troubled by the fact that it is so hard to answer; we lose the motivation for seeking something else—something less "mysterious" than correspondence—with which to identify truth.)

"I accept all that," the pragmatist might say. "Nevertheless, perhaps it is true of the notion of truth (as we find it in practice), that it is identical (in some interesting sense) to ideal warranted assertibility. Shouldn't you therefore allow, at least, the possibility that a Peircean account is the correct one?"

Two points in response to this. The first, an old objection, is that it is very unclear what the notion of the 'ideal limit' might amount to, or even that it is coherent. For example, couldn't actual practice be improved or idealised in several dimensions, not necessarily consistent with one another. In this sense, then, the Peircean pragmatist seems a long way from offering us a concrete proposal. (As Rorty notes, Michael Williams makes a point of this kind.)

The second point—also an old point, for as Putnam famously observes, it is essentially the naturalistic fallacy—concerns the nature of the proposed 'identification' of truth with ideal warranted assertibility. Truth is essentially a normative notion. Its role in making disagreements matter depends on its immediate motivational character. Why should ideal warranted assertibility have this character? If someone tells me that my beliefs are not those of our infinitely refined future enquirers, why should that bother me? My manners are not those of the palace, but so what? In other words, it is hard to see how such an identification could generate the immediate normativity of truth. (It seems more plausible that we begin with truth and define the notion of the ideal limit in terms of it: what makes the limit ideal is that is reaches truth. This doesn't tell us how and why we get into this particular normative circle in the first place.)

I haven't yet mentioned what seems to me to be the most telling argument against the pragmatist identification of truth with warranted assertibility (in Peircean form or otherwise). It often seems to be suggested (by Rorty himself, among others—see the quotes with which I began), that instead of arguing about truth, we could argue about warranted assertibility. This seems to me to miss a crucial point. Without truth, the wheels of argument do not engage; disagreements slide past one another. This is true of disagreements about any matter whatsoever. In particular, it is true of disagreements about warranted assertibility. If we didn't already have truth, in other words, we simply couldn't argue about warranted assertibility. For we could be aware that we have different opinions about what is warrantedly assertible, without that difference of opinion seeming to matter. What makes it matter is the fact that we subscribe to a practice according to which disagreement is an indication of culpable error, on one side or other; in another words, that we already take ourselves to be subject to the norms of truth and falsity.

The crucial point is thus that conversation requires an intolerance of disagreement. This needs to be present already in the background, a pragmatic presupposition of judgement itself. I am not a conversationalist, a judger, at all, unless I am already playing the game to win, in the sense defined by the third norm. Since winning is already implicitly characterised in terms of truth, the idea of a conversational game with some alternative point is incoherent. It is like the idea of a game in which the primary aim is to compete—this idea is incoherent, because the notion of competition already presupposes a different goal. (Here, incidentally, we see the flaw in the pious sentiments of Grantland Rice (1880–1954):

> For when the One Great Scorer comes,
> To write against your name,
> He marks – not that you won or lost –
> But how you played the game.

The One Great Scorer might assign marks on this basis, but we couldn't play the game with such marks as our primary goal; for then it would be a different game altogether.)

There is a connection here with the old objection to relativism, which tries to corner the relativist by asking her whether she takes her own relativistic doctrine to be true, and if so in what sense. The best option for the relativist is to say that she takes the doctrine to be true in the only sense she allows, namely, the relativistic one. When her opponent replies, "Well, in that case you shouldn't be troubled by the fact that I disagree, because you recognise that what is true for me need not be true for you, and vice versa," the relativist has a reply. She can defend the claim that truth is relative to communities, not to individual speakers, and hence that disagreements don't necessarily dissolve in this way.

My pragmatist opponents fare less well against an analogous argument, I think. The basic objection to their position is that in engaging with me in argument about the nature of truth (as about anything else), they reveal that they take themselves to be subject to the norm whose existence they are denying. If they didn't take themselves to be subject to it, they would be in the same boat as the Mo'ans, with no reason to treat the disagreement between us as a cause for concern. They affirm P, I affirm not-P; but by their lights, this should be like the case in which they say 'Yes' and I say 'No', in answer to the question 'Would you like coffee?' (This is what it should be like even if P is of the form 'Q is warrantedly assertible'.) The disagreement simply wouldn't bite. (However, although this argument counts against the (reductive) pragmatist, it is important to note that it doesn't count in favour of a metaphysical realism about truth. Again, the interesting question is why practice is like this, not what truth 'is'.)

6. Metaphysical twists

In effect, I have argued that what matters is that speakers believe that there is a norm of truth, not that there "actually" be such a norm. Isn't this a kind of anti-realism about truth—more precisely, in the current jargon, a form of fictionalism about truth?

In response to this, the first point I want make is ad hominen. I don't think that Rorty can consistently make this objection, because, as he says, he wants to walk away

from realist/antirealist disputes. In other words, he doesn't think that there is an interesting philosophical question as to whether our commitments "mirror" reality.

But can I follow Rorty in walking away from these issues? Or does my 'defence' of truth over justification force me to stay? Perhaps I have to defend realism, to stave off the threat of conversational nihilism—in other words, the threat that if truth does play the role I have claimed in conversational practice, the concession that it is in some sense a fiction would radically undermine that practice, by making it the case that we could no longer consistently feel bound by the relevant norms.

I think I can follow Rorty in walking away. For one thing, nihilism doesn't seem to be a practical threat. In practice, I think we find it impossible to stop caring about truth. (This isn't in itself an argument for realism. The discovery that our biological appetites are not driven by perception of pre-existing properties – the properties of being tasty, sexually attractive, or whatever – does not lessen the pull of those appetites, but no one thinks that that is an argument for realism about the properties concerned.) And even if nihilism were a practical threat, this wouldn't be reason for thinking the claim that truth is a fiction is false (within the game as currently played). It might be a pragmatic reason for not publicising the conclusion too widely, but that is a different matter altogether (at least according to my realist opponents).

More importantly, however, this approach to truth promises to deprive my would-be realist opponents of the conceptual resources they need to make their case. This is because the metaphysical issues tend themselves to be framed in terms of truth, and related notions: theism is said to be in error in virtue of the fact that its claim are not true; that its terms fail to refer. In other words, as I said at the beginning, the issue of the status of truth is enmeshed with the terms of the problem, in a way which is quite uncharacteristic of other cases. For this reason, it is uniquely difficult to formulate a meaningful antirealism or fictionalism about the semantic terms themselves. The right response to this is not to think that we thereby have a transcendental argument for semantic realism. The circularity which undermines antirealism also drains realism of any

interesting content. The right response – as Rorty himself in any case urges – is to be suspicious about the realist/antirealist debate itself.

In other words, not only is my approach deflationary about truth, in the sense that it rejects the idea that there is a substantial metaphysical issue in this case (a metaphysical issue about the 'real' nature of truth – about the truthmakers of claims about truth!). But also, because it is about truth, it positively prevents 'reinflation': it supports and is quite consistent with a general deflationary attitude to issues of realism and antirealism.

7. Conclusion

Rorty has described representationalism – the view that our commitments are answerable to standard beyond ourselves – as the last vestige of a kind of religious deference to external authority. He recommends that we should try to rid ourselves of this representationalist picture, as we have begun to rid ourselves of theism. On this view, realist truth is part of a quasi-religious myth, which we would do better without.

I have agreed that truth is in some sense a myth, or a human creation, but denied that we would necessarily be better off without it. On the contrary, I think, it is a myth that plays an absolutely crucial role in a behavioural pattern – conversation itself – which of great importance to us, as we currently are. It is not clear whether we could coherently be otherwise – whether we could get by without the myth.[6] My claim is simply that we haven't understood truth until we understand its role in the game we currently play.[7]

---

[6] Professor Lloyd's paper at this conference suggested to me that this is very much the same question as whether we could live as thoroughgoing Pyrrhonian sceptics.

[7] I am grateful to Paul Redding and to students in our "Natural and Pragmatism" seminar for many helpful comments on an earlier version.