

When Truth Gives Out

Mark Richard, *When Truth Gives Out*, Oxford UP, 2008, 184pp., \$30.00 (pbk), ISBN 9780199587285.

Reviewed by Paul Boghossian, NYU

2010.05.14

According to Mark Richard, there are important classes of judgment -- genuine, honest-to-goodness claims about the way the world is -- for which (absolute) truth gives out: either because such judgments involve claims that are at best capable only of *relative* truth, or because, more dramatically, they are incapable of truth of any kind, absolute or relative.

Examples of the former sort of judgment include those that involve gradable adjectives ('rich,' 'tall,' 'flat') and those that involve predicates of personal taste ('delicious,' 'tasty,' 'cool'). Examples of the latter sort of judgment include those that involve racial epithets, vague predicates applied to borderline cases, paradoxical sentences, and those that involve predicates of personal taste. The astute reader will have noticed that judgments involving personal taste show up on both lists. This is not because the domain of personal taste is held to include two disjoint sets, one of which consists of truth-apt contents and the other of which consists of non-truth-apt ones. Rather, it is because, on Richard's view, a particular content's belonging to the one list or the other is not a function of its intrinsic properties, but is itself a function of the perspective of the thinkers who judge it. In the case of contents of personal taste, there are perspectives that render them truth-apt and others that render them non-truth-apt. One and the same aesthetic content, then, can be truth-apt relative to one perspective, yet not truth-apt relative to another.



As this brief summary indicates, Richard's views are provocative and intriguing. The book as a whole is smart, stimulating, chock full of dense argument and exhibits a very impressive range.

On the other hand, many of its central claims are underdeveloped, and some seem plagued by serious difficulties. It's not always an easy read. And the individual chapters - one each on racial slurs, vagueness and the semantical paradoxes, emotivism, gradable adjectives, and judgments of personal taste -- for all that they are supposed to illustrate an overarching theme, more closely resemble self-standing papers than they do the nodes of a sustained argument.

Non-Truth-Apt Judgments

The suggestion that there can be genuine judgments that are not truth-apt will be resisted from at least two vantage points.

The first is occupied by the minimalist about truth, who insists that our understanding of the concept of truth is exhausted by our willingness to affirm all instances of the schema

(DS) It is true that p iff p

where 'p' is any judgeable content. Such a minimalist will not understand how it is possible to refuse to predicate truth of a content that one is willing to judge; whenever one is willing to judge that p, one ought to be willing to judge that 'It's true that p,' on pain of failing to understand something constitutive of the notion of truth.

Richard's response to this objection is brief, but I find it persuasive. He says that there are no clearly decisive arguments in favor of minimalism, and so one ought to be open to the claim that there are better reasons for accepting a theory (such as his) that requires minimalism's rejection, than there are for minimalism itself.

The second vantage point is occupied by the thinker who, far from finding truth uninteresting, in the way that the minimalist does, gives it a crucial role in explicating the idea of judgment. To judge something, he will say, is to represent the world as being

a certain way. But a representation of the world (putting vagueness to one side) has to be either correct or incorrect. And how is such correctness to be understood, if not in terms of truth?

Furthermore, it must be possible for beliefs to figure as premises and conclusions of arguments. Arguments, however, must be assessable in terms of their validity; and how else is validity to be understood, if not in terms of truth preservation? Hence, judgments must be the sorts of things that are assessable for truth.

These observations pose a serious challenge to the idea of non-truth-apt judgments. To take such an idea seriously, it is clear, will require some fairly radical reworking of central notions. Before we embark on such a path, we should be convinced that there are central examples of judgments that, intuitively, demand being thought of in this way. Are there such examples?

Richard makes a good *prima facie* case for thinking that there are. As he realizes, what's crucial is to show not just that there are judgments that look to be neither true nor false; rather, what's crucial is to show that there are judgments that seem *correct*, though not by being true, and others that seem *incorrect*, though not by being false. Examples such as these would show that one ought to be able to explain the correctness, or lack thereof, of a judgment, in terms other than its truth or falsity.

Thus, consider the sentence

(1) Prince Charles is a Limey.

We wouldn't want to say that (1) is true, for that would be to endorse the thought expressed; and only a bigot would want to do that. On the other hand, one wouldn't just want to say that (1) is false because that would then commit one to asserting

(2) Prince Charles is not a Limey.

and that would appear to be just as offensive as (1).

Or consider the famous Liar sentence

(L) The sentence L is not true.

(L) cannot be true, for if it were, it would have to be untrue. So, it's not true. But what (L) says is that (L) is not true. So, we can't just say that it's false. So, (L) is neither true nor false.

Something similar could be said to be going on in the case of a vague predicate applied to a borderline case, assuming that the application of a vague predicate to a borderline case results in a claim that is neither true nor false. Given this assumption, and that Jo is a borderline case of baldness, the claim that we make when we utter

(3) Jo is bald.

is not true. But we can't just say that it's *false* that Jo is bald, since it's also not true that Jo is not bald. The sentence

(4) It's not the case that Jo is bald and also not the case that Jo is not bald.

seems to say something correct, but, since it has the form of a contradiction, not, it would appear, by saying something true.

There is, then, a good *prima facie* case for the view that there are non-truth-apt judgments and, along with it, for the affiliated claim that there are judgments that seem *correct* despite not being *true*. But it's hardly a decisive case. With respect to each of these examples, there are alternative ways of thinking about it that Richard cannot be said to have refuted.

So, the question is not whether there is a *prima facie* case for the view, but whether there is an *ultima facie* case for it. And that depends heavily on how satisfactory Richard's non-truth-based alternative is. To the extent to which I have been able to get a handle on it, it seems unpromising.

Richard's Appropriateness Semantics

Richard's basic idea about how to account for these phenomena is to deny that all saying is asserting. In addition to asserting something, he claims, we need to recognize that there is the speech-act of *rejecting* something. These speech acts can be thought of as consisting in, or individuated by, certain *commitments*.

Richard tells us that whereas asserting *p* involves *committing* oneself to the truth of *p*, *rejecting p* is not to be understood as committing oneself to the truth of any proposition. In particular, it is not to be understood as committing oneself to the truth of 'not-*p*,' nor to the truth of the proposition that '*p* is not true' (the 'not's here being understood as expressing the standard truth-functional connective). Rejection is not to be understood as committing oneself to the truth of anything. That is the job of assertion.

So, asserting *p* incurs a commitment to *p*'s being true. And rejecting *p* incurs a commitment to *p*'s being either false or without truth-value. (Somehow this is to be understood as not being equivalent to someone committing himself to *its being true that p is either false or truth-value-less*. For, otherwise, we would be able to assess these commitments in terms of their truth.)

Now, although these commitments can't be assessed in terms of truth, they can be assessed in terms of their *appropriateness*, and it is then natural to say that the positive commitment incurred by asserting *p* is *appropriate* provided the proposition is true, inappropriate otherwise. Correlatively, the negative commitment incurred by rejecting *p* is appropriate provided the proposition is false or without truth-value, inappropriate otherwise.

A *first-order commitment* is a set of such positive and negative commitments. A *second-order commitment* is a commitment to at least one among some class of first-order commitments being *appropriate*.

Richard's suggestion, then, is that we should regard utterance of a sentence intended to be "fact-stating" as utterance that expresses a second-order commitment. Being competent in a language such as English involves knowing what sort of commitment is (conventionally) expressed by serious utterance of its sentences.

Assuming that sentence S expresses the claim p , assertive utterance of S expresses a positive commitment to p and thus expresses the first-order commitment $\{\langle \text{assert}, p \rangle\}$. It also expresses the second-order commitment that it's appropriate to assert p , .

How would one go about rejecting p ? According to Richard, some of our uses of 'not' have precisely the function of expressing rejection rather than expressing the standard truth-functional connective. So there would be a use of

(5) It's not the case that β .

that would express rejection of p rather than the assertion that not- p . Suppose I were to employ that *forced* use of 'not' in uttering (5). Then I would have expressed the first-order commitment $\{\langle \text{reject}, p \rangle\}$, as well as the second-order commitment .

Similar "ambiguities" are postulated for each of the connectives: each connective either expresses a standard truth function or acts as a forced "commitment operator." The word 'and,' in its commitment-expressing guise, "adds up" commitments: if uttering A involves assuming a commitment c , and uttering B involves assuming a commitment c^* , then uttering A and B involves assuming both c and c^* . Similar characterizations are given for the other connectives. Richard gives a systematic way of assigning appropriateness conditions to sentences in which 'not,' 'and,' 'iff,' and so forth, understood as force connectives, are embedded.

Richard then shows how to develop a conception of correct saying in terms of the appropriateness of the commitments expressed and how to explain the validity of inferences involving combinations of forced and unforced (truth-functional) connectives -- in terms of the preservation of apt commitment rather than the preservation of truth.

And so the explanation for why it is OK to utter

(JT) 'Sentence L is not true' is true iff sentence L is not true.

is that it is not an assertion but rather a speech act whose constitutive commitment is that

Either it is apt to assert both of the sentences flanking the 'iff' or to reject both of them.

Since the second of these patterns is appropriate, the sentence is OK to utter.

Understanding Aptness Conditions

My main problem with Richard's picture, I'm afraid, is very elementary. I have not been able to make sense of his notion of rejection and of its constitutive aptness conditions. Rejection, we have been told, is primitive, so we either get it or don't. But its aptness conditions are supposed to be substantively specifiable, because they are needed for determining the aptness conditions for logically complex sentences. In a sign that something has gone seriously wrong, however, I cannot find a coherent and informative way of saying in what the aptness conditions of rejection consist. You couldn't teach rejection's aptness conditions to anyone who didn't already know them.

As noted above, what we are told is that assertion and rejection are (or incur) "commitments," with the former being a commitment to a proposition's being true and the latter a commitment to the proposition's being untrue (either false or truth-valueless). The sentence Richard uses to specify this aptness condition is something like the following:

(Rejection) Rejecting p is appropriate if and only if p is not true.

Do we understand what (Rejection) says? Initially, it might seem as if we do. We understand well enough what it is for a proposition to be untrue -- either false or without truth-value. What could go wrong with stipulating that performing a particular speech act with that proposition is appropriate just in case the proposition has either one of those two properties?

The trouble is that this gloss on (Rejection) involves reading the connectives 'not' and 'if and only if' as truth-functional connectives, and so the serious utterance of (Rejection) as itself an assertion.

But this could not be the intended reading. If it were, it would end up committing itself to the *truth* of

(6) Rejecting p is appropriate if and only if p is not true.

And (6) in turn will be true just in case 'Rejecting p is appropriate' and 'p is not true' have the same truth-value. But that can't be the intended reading, for that would imply that the aptness condition for rejecting p is that it be *assertible* that 'p is not true.' And the whole point was that rejecting p was supposed to be distinct from committing oneself to the assertibility of anything.

So, the 'not' in 'p is not true' cannot be the truth-functional 'not' but must rather be the commitment-indicating or *forced* use of 'not.' But on that reading, (Rejection) would end up saying that

(7) Rejecting p is appropriate if and only if p is not true.

where '*p is not true*' (italicized) expresses not the claim that p is not true but rather the rejection of p.

Now, it's very unclear how we are to understand what is expressed by the rejection of a proposition when it is embedded in a biconditional in this way. Do we read it as expressing an act of rejection? Do we read it as expressing the thought that 'it is appropriate to reject p'? There is a problem either way.

First, let's try to read it as expressing an act of rejection. If we do this, we would then have to read the 'if and only if' as itself a force-indicating operator because, as Richard admits elsewhere, we cannot make sense of force-indicating operators being in the scope of a truth-functional connective, for what could it mean to say something of the form

(8) A if and only if Reject p!

with the 'iff' read truth-functionally? According to Richard, though, utterance of sentences of the form

(9) A iff B.

where the 'iff' is force-indicating, are to be read as expressing one's second-order commitment to its being appropriate either to assert both A and B or to reject both A and B. But what could (Rejection) possibly amount to on this reading? It would look something like this:

(Rejection) It's either appropriate to assert both*

'Rejecting p is appropriate' and 'Reject p!'

or it's appropriate to reject both

'Rejecting p is appropriate' and 'Reject p!'

I have no idea what this is supposed to be saying.

Alternatively, suppose we think of the rejection of p expressed by the RHS of (7) as expressing the commitment that 'it is appropriate to reject p.' But then (7) would end up expressing the trivial thought:

(7) Rejection of p is appropriate iff it is appropriate to reject p.*

Clearly, (7*) can't be used to explain to anyone in what the aptness conditions for rejection consist.

Given all this, it follows that we have no real handle on how Richard's aptness semantics is supposed to work.

The problem here seems to derive from the fact that when one is told that "rejecting p" is appropriate just in case p is phi, one seems to have no choice but to understand this as saying that rejecting p is appropriate just in case p has the *property of being* phi, which then seems to entail that rejecting p is appropriate just in case it's appropriate to assert that p is phi.

Richard on Expressivism

Richard next moves on to extending his commitment semantics to normative sentences, the idea being to show that his picture does a better job of formulating an expressivist semantics than standard accounts, which associate non-truth-apt *attitudes* with normative sentences rather than commitments.

Why would commitments do better than attitudes? Richard's answer is that it is because we can expect the field of commitments -- or the field of their appropriateness conditions -- to be closed under algebraic operations, while the field of attitudes is not so closed. I'll come back to this.

To develop a commitment semantics for normative discourse, Richard suggests looking first at a simple normative sentence:

(S) 'Hunting for sport is good.'

Then, one uses the following procedure: Start by determining what attitude the sentence is a means to expressing. Next, determine what commitment is characteristic of the attitude expressed. And then assign such commitments to the sentences.

(S), we may suppose, expresses *approval* of hunting for sport. But is it so clear that approval involves a distinctive sort of *commitment*? Richard says that approving of x always involves a commitment to x's being worthy of approval. But this is not so clear to me. Can't one simply like something and so approve of it?

Ignoring this issue, what is it to be committed to something's being worth valuing? Richard says that for Jack to commit to x's being worth valuing is for Jack to commit to its being the case that Jack's interests and the facts give him to reason to be favorably

disposed towards x (79). And, therefore, that Jack's commitment will be apt provided that Jack's interests and the facts do give him a reason to be favorably disposed towards x.

If all this were accepted, we could then say that the sentence 'hunting is good' expresses approval of hunting because it is, in virtue of its meaning, a conventional means for expressing the associated commitment. And an expressivist semantic theory would in one way or another tag the sentence with this commitment.

Why is this a better way of doing expressivist semantics than the standard way, which associates *attitudes* with normative sentences? As I noted above, Richard's answer is that the field of commitments -- or the field of their aptness conditions -- can be expected to be closed under algebraic operations, while the field of attitudes cannot.

The aptness conditions of an approval-type commitment are represented by a set Q of interest-world pairs: its complement is just the set of interest-world pairs not in Q. If 'hunting is good' express a commitment apt in just the set of those $\langle i, w \rangle$ such that someone with exactly interests i has, given the facts of w, reason to favor hunting, 'hunting is not good' will express its complement -- which is, of course, the set of those locations "in the space of commitments" in which one's interests do *not* give one reason to value hunting. Something analogous can be said of the commitments expressed by conjunctions and disjunctions of normative sentences, which will be straightforwardly determined by union and intersection of the commitments of their sub-sentences.

Likewise for combinations of descriptive and normative commitments. According to Richard, what unites those who think

(W) If it's bad, Wilhelm won't do it.

is a certain commitment. It is far from clear, he says, that there is a mental *attitude*, in any useful sense of mental attitude, which is necessary and sufficient for having this commitment. But there is no need, he says, if one thinks the primary semantic fact about a logically simple sentence like 'hunting for sport is bad' is that *it* expresses a mental attitude, to think that the same thing must be true of the conditional. An "expressivist

semantics" should stick with the idea that since a sentence's expressing an attitude is a matter of its use making the attitude's characteristic commitment manifest, semantics should assign *commitments*, not attitudes, to sentences.

There is something to Richard's claim that the field of the attitudes isn't closed under algebraic operations. Even if we take it that 'hunting is good' expresses approval of hunting and 'fishing is good' expresses approval of fishing, it's not guaranteed that there is a mental attitude -- and certainly not one of approval -- that corresponds to 'either hunting or fishing is good.' It isn't clear what it would be to approve of either hunting or fishing, but not necessarily one or the other.¹

However, I also don't see just saying with Richard that it's an option for the expressivist to say that there is no mental attitude at all that corresponds to uttering 'either hunting or fishing is good.' After all, one can *think* 'either hunting is good or fishing is good' as an occurrent mental event. This needn't be expressed in language and so can't just be thought of as a speech act. And it needn't just be a matter of incurring certain commitments, since there would clearly be something mental going on.

I think it would be better for Richard's expressivist to say that there always is a corresponding attitude, but that it needn't be one of the familiar ones that we designate through the ordinary mental lexicon, but can only be individuated by its constitutive commitments.

However, Richard's expressivist faces a raft of issues that are not so easily dealt with. First, his view appeals to an unreduced notion of my interests giving me a reason to be favorably disposed towards something. One would have thought that avoidance of such normative facts about practical reason is one of the principal motivations for expressivism.

Second, if when I say 'hunting is good' I express a commitment that's apt just in case my interests give me reason to favor hunting, and when you say 'hunting is not good' you express a commitment that's apt just in case your interests don't give you reasons to favor it, there clearly is no real disagreement between us. But that seems both like the wrong result and one of the Achilles' heels of standard expressivism.

Finally, Richard's expressivist suffers from a problem that will be familiar from our discussion of rejection above. There is a real difficulty stabilizing our conception of commitments from a drift into truth-apt states. Richard warns against succumbing to this temptation:

I expect the reader will be tempted to 'assertize' utterances which involve assuming second-order commitments ... For I expect the reader to reason as follows: If uttering

If hunting is bad, Wilhelm won't do it.

has aptness conditions and uttering it is committing to those conditions being in place -- to their obtaining -- why not just look upon such utterance as an assertion ... of something that is true iff such conditions obtain? The advantage, it will be thought, is that we achieve a uniform story about (more) utterances, taking (more of) them to be assertions, and thus not needing (quite such) a novel, complicated story involving armies of novel speech acts. (77)

His answer doesn't meet the worry head-on, but says rather that if we insisted in thinking that way we would no longer be able to avail ourselves of his satisfying treatment of the paradoxes and vagueness, as previously outlined.

But we already know this line of thought can't be correct. I can't be asserting of the claims made by the sentences flanking the 'iff' in

'Jo is bald' is true iff Jo is bald

(Jo is a borderline case of baldness) that they are either both assertible ... or are both not ... assertible. At least I can't be asserting this and saying something true. There's nothing I might assert with the biconditional that's true... . I take it that because the story I've told about the liar, vagueness and denial is so satisfying -- it is so obviously correct -- we know this to be so. (77)

Given the problems previously outlined for the notion of rejection, as Richard characterizes it, it's hard to share his confidence. Indeed, it's hard to see how these new normative commitments could avoid the drift towards truth-aptness that the notion of rejection seemed to face.

The commitment expressed by (W) is apt iff either (A) one's interests and the world don't provide one with reason to disapprove of hunting or (B) Wilhelm won't take up hunting.

It is easy to understand this as saying: The commitment expressed is apt iff either it's true that (A) or it's true that (B). But that can't be the intended reading because that would make the commitment expressed indistinguishable from a judgment to the effect that it's either true that (A) or true that (B).

But what other intelligible reading is there?

Richard on Relative Truth

Richard's next big topic is that of relativism and relative truth. According to him, while some claims are incapable of truth, others that are truth-apt are capable only of *relative* truth. In particular, he claims, any sentence involving a gradable adjective must be given such a relativistic treatment.

The idea that certain classes of contents must be seen as having only relative truth values and not absolute ones has become extraordinarily popular in recent analytic philosophy. The trend has been fueled to a considerable extent by a younger generation of philosophers – John MacFarlane and Max Kölbel, for example -- but senior figures like Richard have also been important contributors.²

Richard's main example runs as follows. Mary has just won a million dollar lottery and Didi, impressed with this, says:

(D) Mary is rich.

Naomi, for whom a million dollars is not that much, remarks in a conversation disjoint from Didi's,

(N) Mary is not rich at all.

It is well understood that sentences involving gradable adjectives can express different propositions in different contexts of use, depending on the comparison class that is operative within the context. Thus, what Didi may have meant is that Mary is rich for a resident of Detroit while Naomi, speaking among her rich pals in New York may have meant that Mary is not rich for a New Yorker. All this, of course, can and should be understood along standard Contextualist lines.

According to Richard, though, gradable adjectives are contextual in a further sense: even if two speakers associate exactly the same comparison class with 'rich,' they may nevertheless use it to pick out different properties in different contexts of use, by virtue of having somewhat different ideas about where the cut-off point is between the rich and the non-rich.

To focus on that aspect of things, let's legislate that in the case of Didi and Naomi the relevant comparison classes in the two conversations is exactly the same (New Yorkers, let's say) and also that there is no difference between the two conversations so far as the point of assessing people as rich goes -- each of the women is assessing people as rich or not and then assessing whether the rich ones deserve their wealth.

It is part of the example, then, that Didi and Naomi are operating with somewhat different standards (cut-off points) for being rich, even given the fixed comparison class. (Maybe Didi thinks the cut-off is around a quarter million dollars and Naomi thinks it's more like two million dollars.) Their interests, values and backgrounds make them draw the line between the rich and non-rich in somewhat different places.

It is also part of the example that this is known, or knowable, to all concerned. If Didi and Naomi were to be placed in the same conversational context, they could be presumed to know that each of them is operating with a somewhat different standard for being rich. Finally, their judgments, (D) and (N), are true relative to their respective standards and each of them is justified in thinking that they are.

Relativism and Faultless Disagreement

Now, how do we get from these materials an argument for relativism as opposed to Contextualism? As we saw above, what we need is some reason to think that Didi and Naomi genuinely disagree with one another but also that their judgments are capable *at most* of relative truth.

A classic line of thought, applied to this case, would go something like this (Max Kölbel may have been the first to lay it out with some care).³

We start out by claiming that it doesn't merely *seem* that Didi and Naomi disagree with one another, but that they genuinely do so. Moreover, their disagreement is a substantive, factual one as to whether Mary is rich, not a metalinguistic one, about what 'rich' means or how any vagueness in 'rich' is to be resolved. This entails that Didi affirms and Naomi denies one and the same content -- viz., *Mary is rich*.

(Disagreement) Didi and Naomi disagree with one another in the sense that one affirms and the other denies the content Mary is rich.

This in effect rules out a Contextualist view of the case, according to which Didi and Naomi are not really disagreeing because what they mean by their utterances is, respectively,

(D1) According to standards S(D), Mary is rich.

and

(N1) According to standards S(N), Mary is not rich.

Second, this disagreement is *faultless* in the sense that neither Didi nor Naomi could be said to have committed a mistake in believing what they do. Here is Kölbel's characterization of faultless disagreement:

(Faultless Disagreement) A faultless disagreement is a situation where there is a thinker A, a thinker B, and a proposition (content of judgement) p such that:

- (a) A believes (judges) p and B believes (judges) not-p,*
- (b) Neither A nor B has made a mistake (is at fault).⁴*

If we put disagreement and faultlessness together, we seem to get a neat argument for relativism. Didi and Naomi disagree with one another: in particular, one of them asserts p and the other not-p. Yet neither of them is guilty of a mistake. If neither of them has made a mistake, then both of their beliefs are true ("mistake" here is not purely epistemic). If their beliefs were capable of absolute truth, we would have to allow that there are true contradictions. But we don't want to allow that. Hence, their beliefs must be at most capable of *relative* truth. This allows both Didi and Naomi to have correct (faultless) beliefs, even as they disagree with one another.

Now, crucially, Richard *rejects* this familiar and influential line of thought. His reasons are important. He says that a commitment to the existence of faultless disagreement cannot be sustained in this case

simply because when one is willing to ascribe truth or falsity to a particular claim p , one treats p and the claim that p is true as equivalent: within a perspective, truth is 'disquotational'. Suppose I think that Beaufort is a better cheese than Tome, and you think the reverse. Suppose (for reductio) that each of our thoughts is valid -- mine is true from my perspective, yours is from yours. Then not only can I (validly) say that Beaufort is better than Tome, I can (validly) say that it's true that Beaufort is better than Tome. And of course if you think Tome is better than Beaufort and not vice versa I can also (validly) say that you think that it's not the case that Beaufort is better than Tome. So I can (validly) say that it's true that Beaufort is better than Tome though you think Beaufort isn't better than Tome. From which it surely follows that you're mistaken -- after all, if you have a false belief, you are mistaken about something. This line of reasoning is sound no matter what the object of dispute. So it is just wrong to think that if my view is valid -- true relative to my perspective -- and your contradictory view is valid -- true, that is, relative to yours -- then our disagreement is 'faultless'. Faultless disagreement is possible -- but such disagreement is not one to be evaluated in terms of truth. (132)

Let us call this the Argument from (Perspectival) Immersion. It will prove useful to lay it out with some care. Suppose that D judges p , N judges not- p and that both judgments are valid (i.e., true relative to the individual perspectives from which they are judged). According to Richard, the alethic relativist about the content p must hold all of the following:

The Argument from (Perspectival) Immersion:

(10) *The content (p) is at best relatively true. (Alethic Relativism)*

(11) *If D judges validly that p, it will also be valid for D to judge that It's true that p. (Truth is Disquotational within a perspective)*

(12) *If D judges that It's true that p then D must, on pain of incoherence, judge that It's false that not-p. (We are not in the realm of the gappy.)*

(13) *If D judges that It's false that not-p, then D must, on pain of incoherence, judge that anyone who judges not-p (e.g., N) is making a mistake.*

Therefore,

(14) *D must judge that N is making a mistake and so cannot regard the disagreement with N as faultless.*

Therefore,

(15) *The disagreement between D and N is not faultless.*

As we shall see, Richard allows that there *can be* faultless disagreement, though when there is such disagreement the judgments in question can't be evaluated in terms of their truth at all, not even in terms of their *relative* truth.

Relativism without Faultless Disagreement

If Richard does not motivate relativism about gradable adjectives using a Kölbel-style argument from faultless disagreement, how does he do it?

Richard's distinctive line here is to arrive at relativism through a combination of what he calls Accommodation and Negotiation.

Accommodation, of course, is David Lewis's word for the way in which the extension of a word can shift in a conversational setting so as to make sentences in which it is used in that conversation true, provided no one objects to the use in question.

Is France hexagonal? Well, relative to certain low standards it counts as hexagonal. In a given conversational context, one's interlocutors might accommodate those low standards and so, in that context, accept 'France is hexagonal' as true. Lewis's claim was that that sufficed to make it *true* in that context. In another conversational context,

dominated by other interests, the relevant participants might not accommodate those low standards and so, in that setting, would not accept 'France is hexagonal' as true. Lewis's claim was that that would suffice to make it false in that context.

Even if we accepted all this, there would be no argument for relativism because all this could be perfectly well accommodated by the Contextualist, as Lewis himself was clear about. Speaking of Peter Unger's provocation to the effect that nothing is flat, he writes:

The right response to Unger, I suggest, is that he is changing the score on you. When he says that the desk is flatter than the pavement, what he says is acceptable only under raised standards of precision. Under the original standards the bumps on the pavement were too small to be relevant either to the question whether the pavement is flat or to the question whether the pavement is flatter than the desk. Since what he says requires raised standards, the standards accommodatingly rise. Then it is no longer true enough that the pavement is flat. That does not alter the fact that it was true enough in the original context. "The desk is flatter than the pavement" said under raised standards does not contradict "The pavement is flat" said under unraised standards, any more than "It is morning" said in the morning contradicts "It is afternoon" said in the afternoon.⁵

But Richard also thinks it's important to emphasize that 'rich' is subject not just to accommodation but also to 'contextual negotiation,' (an idea he attributes to Julius Moravscik). When speakers differ over how 'rich' is to be applied to cases, he says, they can and often do attempt to reach a consensus as to how it is to be applied, via examples, argument, mutually agreeable stipulations and so on. And, as previously emphasized, such disagreement can often be factual and not just metalinguistic.

So here we have a combination not of Disagreement and Faultlessness but of Disagreement and Accommodation. And we can trace an alternate route from this latter combination to relativism. Because there is accommodation, speakers are able to impose different extensions on the word 'rich' in different conversational contexts. However, Didi and Naomi, made aware of their divergent verdicts about the extension of 'rich,'

may subject their uses of 'Mary is rich' to contextual negotiation. The fact that there is something substantive to negotiate about suggests that Didi and Naomi genuinely disagree. Disagreement implies that there is some content that Didi affirms and Naomi denies. Accommodation implies that one and the same content can be accommodated within one conversational context, and so be true in that context, and yet fail to be accommodated within another conversational context, and so fail to be true in that context -- all this while holding the world fixed. So, once again, relativism.

There are two weaknesses in this argument, it seems to me. First, there really aren't very good ways of deciding when a dispute is factual and when it is in some broad sense verbal. Richard's case depends heavily on supposing that there are and that this particular case falls squarely on the factual side.

Second, even if we could somehow make it seem plausible that the dispute in this case is a genuine disagreement, it's not clear how it can help Richard, because, as we shall see, even his relativistic framework has great difficulty making it come out to be a genuine disagreement.

Two Problems for Alethic Relativism, Given the Argument from Immersion

But before looking at that, I want to describe two problems for Richard's view, assuming it to be well motivated by the combination of Accommodation and Disagreement.

To see what these problems are, let's suppose that I am a Richard-style relativist about the content *Mary is rich* and also a committed thinker, like Didi, who judges that Mary is rich. If I endorse the Argument from Immersion, I will have to conclude that on the question whether Mary is rich -- about which I am supposed to be a relativist -- there is only *one* correct thing to think, namely that Mary is rich. That would seem to be implied by the fact that I must take any conflicting view to be, first, *eo ipso* false and, second, *eo ipso* mistaken.

To many this view will seem to be a relativism in name only. For what most people think is that to be a relativist about a given question is to allow that there can be more than one correct answer to it. To be a moral relativist, for example, on this traditional view, is to hold that, on a given moral question, there can be more than one correct answer, these correct answers not being necessarily consistent with one another. Surely, they

might say, it's odd to claim that one could count as a relativist on the question whether one ought to place one's elbows on the table while eating, while maintaining there to be exactly one correct answer, applicable to anyone, anywhere, no matter what their backgrounds, standards, interests and cultural setting might be. Given this traditional understanding of relativism, relativism cannot be divorced from faultless disagreement because the latter is constitutive of the former.

On this traditional view, then, we should say that *alethic relativism*, in the technical semantical sense that Richard is working with, doesn't *in and of itself* amount to a relativism of a recognizable and intuitive kind, since it is possible to combine its technical framework concerning the dependence of truth on indices with the Argument from Immersion to yield a highly *intolerant*, objectivist view of the subject matter at issue.

To put the matter succinctly, at the core of relativism, intuitively understood, is a *normative* thesis concerning tolerance of opposing views. But to stake out some claims about the dependence of truth on indices is not yet to have made any kind of normative claim. It is only once we have said what the normative significance is of having contents that are capable only of relative truth-values that we can assess whether what we have is relativism of a recognizable and interesting kind -- or something else.

OK. So perhaps Richard's alethic relativism isn't a relativism of a recognizable and interesting kind. Perhaps its interest lies elsewhere.

The problem that we are focusing on at the moment, however, has the potential not merely to raise a question about the relativistic *credentials* of alethic relativism, but to expose a fundamental instability in the view itself. That is what I now want to turn to.

The content *Mary is rich*, I say, cannot be true or false simpliciter. It can only be true or false relative to a standard. I know that my belief that Mary is rich is true relative to my standard, and I know that Naomi's belief that she is not rich is true relative to hers. Still, I think that Naomi's belief is mistaken and I will attempt to get her to change her mind.

What can be my attitude towards Naomi's *standards*? Well, obviously, I must regard them as mistaken. After all, it's part of the example that Naomi is reasoning impeccably from her standards. Since I regard her judgment as mistaken, I must regard her

standards as mistaken. (If she uses a cut-off of 2 million for being rich, and I use 250K I must regard her cut-off as mistaken.) Further, I'm trying to get her to change her mind -- why would I do that if I regarded her standards to be just as correct as mine?

So in the presence of the Argument from Immersion, I have no choice but to regard Naomi's standards as mistaken. Indeed, I have no choice but to regard any standard that yields verdicts incompatible with mine as mistaken. Holding the details of the case fixed, I am committed to there being a uniquely correct standard by which the question of Mary's wealth is to be judged (and obviously I am in possession of it).

Now, how can I cogently say that the truth-value of 'Mary is rich' (p) varies across our perspectives? If, as I am now insisting, there is a *best perspective* on p (mine), shouldn't we say that p *has* an absolute truth-value, and not merely relative ones, namely the truth-value that it has in the 'best perspective'? If the uniquely correct standard for deciding the question about Mary is a cut-off of 250K, then the claim that Mary is rich would appear to be true just in case she has 250K or more, false otherwise.

So, an intolerance about the opposing standards threatens to destabilize the thesis central to Richard's view, that p has only relative truth-values and no absolute ones.

Richard spends a lot of time trying to explain how it need not be disingenuous of me to ask you to abandon a view that is 'true for you' in favor of one that is (currently) 'false for you.' For, he says, to ask you to do that is, in most cases, just to ask you to examine your values and priorities and to get you to change them, even if there are no 'perspective-independent' considerations that I can offer you.

Well, it's no easy task to say how getting someone to change his or her basic, underived values can be a rational process and I don't think that what he offers us is all that helpful. But it's also not really the crucial point. The crucial point is that so long as the view embeds a commitment, for any given dispute about p, to there being a *best perspective* by which that dispute ought to be decided, it's hard to square that with the idea that p only has perspective-dependent truth-values and no absolute ones.

Let me sum up this horn of the dilemma. If the alethic relativist endorses the Argument from Immersion, then not only can his view not be effectively distinguished from a view that is anti-relativist/objectivist in spirit, but its central claim, that p only has relative

truth-values, appears to be destabilized.

Disagreement?

All of this seems to suggest that the relativist would be well-advised to reject the Argument from Immersion. I don't have the space to do so here, but I have argued elsewhere that there is only one semi-plausible way to do so and that is to reject (13) -- the claim that a false belief is *eo ipso* mistaken. I don't myself believe that such a counterintuitive view can ultimately be sustained, but if there is a promising way out for the relativist, it lies along that path.⁶

But even if that could be done, it seems to me that the relativist is in trouble in another way.

As the reader will recall, the relativist view is motivated by the claim that the dispute between Didi and Naomi is a genuine factual disagreement. It is that alleged fact that scotches the prospects for a Contextualist account of what is going on between them. At the end of the day, though, it seems clear that not even the alethic relativist can respect that alleged fact, thus surrendering whatever advantage her view was supposed to have over the Contextualist view.

Richard tries to explain why his picture preserves genuine disagreement between Didi and Naomi. 'Rich,' he says, is to be associated with a context-invariant *notion* (modeled on Kaplan's notion of a character) that takes a context as input and yields a property as output. In different contexts, 'rich' will pick out different properties, even given a fixed comparison class. The objects of assertion, belief and the other attitudes are made up out of *notions*, not properties. To believe that Mary is rich is not to attribute the property of being rich to Mary, but to apply the *notion* of being rich to her.

Contents made up of notions constitute r-claims, contents that are at best relatively true. In the case of the r-claim that Mary is rich, its *completion* in my context is the proposition that Mary has P, where P is the property 'is rich' expresses when I use it. R-claims p and q are *inconsistent* provided that for any context c the *completions* of p and q are incompatible (that is, there is no world in which both completions are true). Didi and Naomi count as disagreeing by virtue of believing inconsistent r-claims.

But this can't be an adequate account of disagreement (never mind the big question whether it could be an adequate account of the objects of belief). Suppose we found it convenient to give a relativistic account of "It's raining." On such a view, we don't regard an utterance of "It's raining" as elliptical for some longer sentence specifying the location at which the utterance is made but treat it rather as already expressing a complete assertable or thinkable content whose truth value varies with the context of utterance but which makes no explicit reference to that context. Then, on Richard's view of disagreement, we would have to say that if I said "It's raining" in NYC and you said "It's not raining" in Chicago, that we would thereby be disagreeing about the weather, even though what each of us said was justified and true relative to our respective locations. But, surely, that's absurd.

So, even if we were to grant that (13) could be cogently rejected, that still wouldn't get us disagreement. For all that the relativist is able to offer us, then, we might just as well have stuck with Contextualism.

Faultless Disagreement without Relative Truth

Richard's concluding chapter is about judgments of taste (e.g., 'this is cool,' 'this is delicious'). Here he holds that one and the same judgment can sometimes be the object of a faultless disagreement and sometimes not:

when we disagree about matters of taste, it is often correct to say that though we disagree neither of us is making an error. From this it follows that our beliefs are not even relatively true or false. But it is also the case that people may disagree about a matter of taste -- they may have the same disagreement as we have -- and one of them may validly judge the other is mistaken ... And from this it follows that the disputed claim enjoys at least a relative truth. That last paragraph might sound a tad inconsistent. It holds out the possibility of there being a claim (over which you and I disagree) that isn't even relatively true or false, although it is identical with a claim (over which some other people disagree) that is relatively true or false. But there's no inconsistency here, given that not only may the assignment of a truth-value be 'perspective-relative,' but whether a claim is truth-apt may also be 'perspective-relative.' ... the right thing to say about matters of taste seems to be that not only is it a matter of taste whether duck-liver mousse is delicious, it is (so to speak) a matter of taste whether the claim that it's delicious is true or false [by which he means truth-evaluable]. (125-6)

Now, as we have just seen, an alethic relativist is not going to be able to avail himself of a classical Frege/Russell view of a judgeable content. Since he wants the truth-value of a content to vary not only with the world relative to which it is assessed but also with the judger's (variable) standard of taste, he cannot be operating with a classical view of propositional content.

The relativist, though, must somehow think that the object of judgment is some character-like thing, since all that's varying from context of utterance to context of utterance is the truth-value.

Now, this commitment is already problematic enough. It's as though we are able to judge the truth of the content "It's raining," but without knowing which place was being referred to, not even under the description "here."

But on Richard's view of taste judgments, things get even weirder. For now we can no longer stably regard a content about a matter of taste as a function from a context to a truth-condition. Rather, sometimes when I judge such a content, it is such a function, and at other times, if it is involved in a disagreement that the interlocutors regard as flawless, it is not such a function. Yet it is one and the same content.

Indeed, it seems as though we can get the following situation: I judge that p, you judge that not-p. I regard the disagreement as flawless, so my p is not truth-apt, but you regard it as lawful and so your p is truth-apt.

How can this disagreement turn out to involve one and the same p? That, however, would appear to be the minimum that's required for this to count as a real disagreement.

Conclusion

As can be seen from the length of this review, I found Richard's book hugely stimulating. Readers will benefit from thinking hard about his various ideas, even if they end up not agreeing with many of them.⁷

¹ See Richard's review of Mark Schroeder's *Being For*, forthcoming in *Philosophical Review*.

² Max Kölbel, *Truth Without Objectivity*, (London: Routledge, 2001); John MacFarlane, "Making Sense of Relative Truth," *Proceedings of the Aristotelian Society* 105 (2005), 321-39.

³ Kölbel, "Faultless Disagreement," *Proceedings of the Aristotelian Society*, 2003, 104 (1): 53-73.

⁴ Kölbel, *ibid.*

⁵ Lewis, "Scorekeeping in a Language Game," reprinted in his *Philosophical Papers I*, (Cambridge: Cambridge UP, 1983), 245-6.

⁶ For discussion, see my "Three Kinds of Relativism," forthcoming in S. Hales (ed.) *The Blackwell Companion to Relativism*.

⁷ I am grateful to Mark Richard, Seth Yalcin, Crispin Wright, Paul Horwich, John MacFarlane and Stephen Schiffer, as well as the participants in a meeting of the Project on Disagreement at the New York Institute of Philosophy, for comments on some of the material that went into this review.