

Richard on Truth and Commitment

John MacFarlane
Department of Philosophy
UC Berkeley
jgm@berkeley.edu

October 19, 2010

The twentieth century has left us with abundant tools for theorizing about the meanings of sentences that can be used to represent states of affairs — sentences that are “truth apt.” But what do we do, Mark Richard wonders, when “truth gives out”? How should we think about the aspects of our thought, talk and reasoning that are not straightforwardly representational: pejoratives, for example, or taste-related evaluatives like “cool”? In *When Truth Gives Out*, Richard develops two theoretical tools for dealing with such discourse. One is a form of expressivism, which I’ll call “commitment semantics.” The other is a form of truth relativism. In what follows, I’ll discuss both of these tools, and the uneasy relation between them.

Commitment semantics

Commitment semantics starts from the idea that many apparently assertoric utterances are best understood not as assertions (commitments to the truth of a proposition), but as other kinds of commitments. Saying that a red leather cape is cool, for example, is committing oneself to valuing the cape in a certain way. And saying that the Liar is not true is undertaking a commitment that is apt just in case asserting that the Liar *is* true would not be apt.¹

The idea, thus crudely put, has much in common with emotivism in ethics—a position that many philosophers take to have been decisively refuted by considerations due to Frege and Geach. The basic objection is that, if we explain the meaning of

¹Most of the time Richard uses the word “appropriate” instead of “apt” (and sometimes he even uses “fulfilled”), but I think “apt” is most apt for his purposes. He wants to say that a commitment to the truth of *P* is “apt” or “appropriate” if *P* turns out to be true, even if the person undertaking the commitment did not have good evidence for *P*, and committed herself recklessly. In such a case, we would not ordinarily say that the commitment was “appropriate,” but we might call it “apt.”

(1) That cape is cool!

by explaining what kind of speech act it is normally used to perform, we are left without any account of its meaning in embedded contexts, such as

(2) If that cape is cool, then I will buy one for my brother.

Nor do we have any account of why the inference from (1) and (2) to

(3) I will buy one for my brother.

is a good one.

This is the problem commitment semantics is designed to solve, by giving us an elegant way to compute the commitments conventionally undertaken by utterances of complex sentences as a function of the commitments conventionally undertaken by utterances of their simpler constituents. Instead of recursively defining truth conditions, as we do in truth conditional semantics, we recursively define commitments (or rather, as we'll see, aptness conditions for commitments).²

Here's how it works.³ First, we can classify simple commitments by *dimension* and *polarity*. Assertion and rejection (discussed in Chapter 2) are commitments in the dimension of *truth*, but they differ in polarity. Assertion is a *positive* commitment to the truth of a proposition, while rejection is a *negative* commitment to its truth—a commitment that is apt iff the positive commitment is not. We can specify any *simple commitment* by specifying its polarity, its dimension, and its object—in the case of assertion and rejection, a proposition; in the case of valuing, an object or state of affairs. So, for example,

(4) +Truth(that snow is white)

(5) -Truth(that snow is white)

(6) +Valuing(that cape)

(7) -Valuing(that cape)

Each simple commitment c has an *inverse* \bar{c} —the commitment that differs from it only in polarity. In the examples above, (4) and (5) are inverses of each other, as are (6) and (7). A simple commitment is apt just in case its inverse is not apt.

²Though Richard doesn't discuss the connection, the idea has some similarities to Simon Blackburn's (2001) use of semantic tableaux to give a systematic accounts of the commitments undertaken by the use of complex sentences. Richard's approach differs significantly from Blackburn's, though, in ways that make it easier to defend.

³In what follows, I present Richard's idea in my own way, which I think makes it perspicuous, though some of the distinctions I make are only implicit in Richard's book.

A *first-order* commitment (FOC) is a set of simple commitments considered conjunctively. An FOC is, essentially, a joint undertaking of all of its members, and is apt iff all of its members are.

A *second-order* commitment (SOC) is a set of FOCs considered disjunctively. An SOC is apt iff at least one of its members is. Any commitment intuitively undertaken in a speech act can be represented as a SOC. In the simplest cases, the SOC might contain just one FOC, which in turn consists of just one simple commitment. So, for example, in asserting that snow is white, one undertakes the SOC

$$(8) \{\{+\text{Truth}(\text{that snow is white})\}\}$$

But more complex cases are possible. So, in saying that either snow is white, or snow is yellow and is not cool, one undertakes the SOC

$$(9) \{\{+\text{Truth}(\text{that snow is white})\}, \{+\text{Truth}(\text{that snow is yellow}), -\text{Valuing}(\text{snow})\}\}$$

How do we compute these complex commitments compositionally? Define the conjunction, disjunction, and inverse operations on SOC's as follows:

$$(10) \text{Conj}(A, B) = \{x \cup y \mid x \in A \ \& \ y \in B\}$$

(This is a commitment that is apt iff both A and B are apt.)

$$(11) \text{Dis}(A, B) = A \cup B$$

(This is a commitment that is apt iff either A or B is apt.)

$$(12) \text{Inv}(A) = \{\{\bar{x} \mid x \in C\} \mid C \subseteq \bigcup A \ \& \ \forall B \in A, |B \cap C| = 1\}$$

(This is a commitment that is apt iff A is not apt. Since A is apt just in case any of its members is apt, $\text{Inv}(A)$ is the set of minimal commitments that exclude the aptness of any member of A .)

These operations give us meanings for conjunctions, disjunctions, and negation of arbitrary “forced” sentences, and allow us to compute the commitment conventionally undertaken by any arbitrarily complex “Boolean” combination of simpler sentences.⁴ We can think about this as a theory of semantic competence, since “[o]ne undertands such a language only if one can tell, for any of its sentences, what sort of speech act it is a vehicle for—only if one knows. . . what second-order commitments one incurs by uttering it” (67).

Richard defines validity for this semantics not as the preservation of commitments, but as the preservation of *aptness* of commitments:

⁴I put “Boolean” in scare quotes because we don’t have the structure of a Boolean algebra here; for example, $\text{Conj}(A, \text{Inv}(A)) \neq \text{Conj}(B, \text{Inv}(B))$.

validity is simply the preservation of the aptness of commitment. that is, an argument is valid provided that whenever the commitments associated with all of its premises are appropriate, so is the commitment associated with its conclusion (68).⁵

Given this semantics and account of consequence, Richard can meet Geach's challenge to emotivism. He can explain how sentences whose meanings are understood non-truth-conditionally contribute systematically to the meanings of sentences containing them, and he can explain the validity of arguments in which they figure.

What are these commitments?

The idea that understanding a sentence is a matter of knowing what commitments one would conventionally undertake in uttering it is an appealing one. But does Richard's semantics really tell us what these commitments are?

Let's start with the simplest, most basic commitments. What commitment do I undertake when I say

(13) It rained yesterday, or

(14) That cape is cool!, or

(15) You may play outside.

Richard's official answers seem to be:

(16) a commitment to the truth of the proposition that it rained yesterday (50),

(17) a commitment to *valuing* that cape in a distinctive way (150), and

(18) a commitment to there being some satisfactory situation in which the proposition that you play outside is true (84).

The problem is that it isn't clear what a commitment to the truth of a proposition, or to the existence of a situation, involves. Commitments are presumably

⁵Elsewhere Richard says that "To say that S *c*-entails A is, roughly, to say that anyone who seriously utters all of the members of S , thereby incurring the commitments associated with those sentences, incurs also the commitment associated with A " (63). "Roughly" indeed; the notion expressed here is not at all the same as preservation of aptness of commitment. Take the inference " $A \vee B, \neg A, \text{ therefore } B.$ " Here we have preservation of aptness of commitment, in the sense that if the commitments undertaken in the premises are apt, so is the commitment undertaken in the conclusion. But we don't have preservation of commitment, since the conclusion may introduce a commitment entirely disparate from the commitments undertaken in the premises.

not just things that can be apt or inapt, but things that one can *violate* or *fulfill*. So, for any commitment, we can ask what sort of action (or inaction) would count as fulfilling or violating the commitment. One fulfills commitment (17), for example, if one properly values the cape, and one violates (17) if one fails to do so. That I understand. But what counts as fulfilling or violating the commitment (16)? Not an act of *making* it true that it rained yesterday—nobody can do that, and even if one could, a commitment to do that would be more like a promise than an assertion. The same questions can be asked about (18). It seems to me that Richard owes an answer to these questions, if he is to be entitled to talk of commitments here.

It is interesting that when it becomes useful for Richard to talk of *fulfillment* of commitments, he says that someone fulfills a first-order commitment to the truth of p if she *asserts* p (60). But it would be very odd, not to mention regress-engendering, to say that the commitment one undertakes in asserting p is a commitment to asserting p .

It is even more obscure what Richard's second-order commitments come to. We are told that the second-order commitment

(19) $\{A, B\}$

is a commitment that is apt just in case the commitment undertaken by A is apt or the commitment undertaken by B is apt. At best, that's a very indirect way of picking out the commitment: instead of being told what would count as fulfilling or violating it, we are given its aptness conditions. (One might know that punting is the move in American football that is apt when the offensive team is facing their final down and is too far from the goal line to score a touchdown, without having any idea how to punt or recognize a punt.) Worse yet, settling the aptness conditions of a commitment does not seem sufficient, in general, for settling what commitment it is, because distinct commitments can have the same aptness conditions.

We can illustrate this point just by thinking about assertions—positive commitments to the truth of a proposition. On Richard's view, an assertion is apt just in case its content is true. Let p and q be two ordinary, truth-apt propositions, and consider the commitment undertaken by an assertion of

(20) $p \vee q$,

namely

(21) $\{\{+\text{Truth}(p \vee q)\}\}$.

This commitment is apt just in case $p \vee q$ is true, so it is apt just in case p is true or q is true. But (22) and (23), which ought to count as different commitments on Richard's theory, are apt in exactly the same cases.⁶

(22) $\{\{+\text{Truth}(\neg(\neg p \wedge \neg q))\}\}$

(23) $\{\{+\text{Truth}(p)\}, \{+\text{Truth}(q)\}\}$

The problem is even worse if we consider

(24) $\{\{+\text{Truth}(p \wedge \neg p)\}\},$

a commitment that is *never* apt, and thus, I assume, has the same aptness conditions as

(25) $\{\{+\text{Valuing}(\text{kicking a homeless person in the gutter})\}\}$

So, when Richard tells us that a second-order commitment is "the" commitment that is apt just in case one of its members is apt, he has not singled out one commitment, but at best indicated a set of commitments.⁷ Nor has he given us any proof that the set is nonempty, for example by directly describing one commitment that has these aptness conditions.

It might seem tempting to say, at this point, that each commitment is a commitment to its aptness conditions' *obtaining*. That looks like a recipe for turning a set of aptness conditions into a direct specification of a commitment. But it would make every commitment vulnerable to the kind of question we posed above for (16) and (18): what would count as fulfilling or violating a commitment to a state of affairs? Worse, it would make every second-order commitment equivalent to an assertion (the assertion that its aptness conditions obtain).

At one point Richard seems to see at least part of the difficulty:

We have a recipe for determining, once we know the conditions under which it is appropriate to do whatever it is that one does in uttering the sentence A , the conditions under which it's appropriate to do whatever it is that one does when one utters *not* A , A or B , and A and B . But that doesn't quite tell us how to get from a specification of what one does, in uttering A and uttering B , to a specification of what one *does* in uttering *not* A , A or B , or A and B . (74)

⁶Note that Richard takes propositions to be structured, so the proposition that $\neg(\neg p \wedge \neg q)$ is different from the proposition that $p \vee q$.

⁷When he first introduces force disjunction, Richard says (60) that for every set of second-order commitments, there is a *unique* second-order commitment that one fulfills just in case one fulfills some of its members. But no proof is offered.

But his response is really just a refusal to respond:

So long as the commitments have appropriateness conditions to begin with, it will make perfect sense not simply to incur them, but to disavow, disjoin, and conjoin them. Again, deeper explanation of what one *does* in uttering a force negation (disjunction, conjunction) when speaking the language is neither possible, necessary, nor desirable. (75)

What's wrong with relativism?

In Chapter 4 Richard gives an able defense of a form of truth relativism, arguing that a relativist semantics is needed to make sense of the distinctive way in which terms like "rich" and "flat," as well as philosophically interesting terms like "knows," are contextually sensitive. Here's the basic setup. Sitting in a small New York restaurant, I happen to overhear two conversations about Mary, an acquaintance who just won a million-dollar lottery:

(26) Didi (in conversation 1): Mary is rich.

(27) Naomi (in conversation 2): Mary is not rich at all.

I presume that Didi and Naomi are both using rich to mean "rich for a Manhattanite." It is then appropriate for me to say:

(28) Didi and Naomi disagree about whether Mary is rich.

If there is any doubt about this, suppose Didi and Naomi are told of what the other said, and begin to argue. Surely they disagree about whether Mary is rich. Since each takes herself to be defending the claim she made before the argument began, Richard observes, it seems plausible that they disagreed already in making these claims.

However, it *also* seems plausible that

(29) Didi's claim (26) is true in the context of conversation 1.

(30) Naomi's claim (27) is true in the context of conversation 2.

When we use the word "rich," we do not assume that there is some neutral, non-contextual fact about where the cutoff point between rich and non-rich New Yorkers lies. We let the cutoff point shift contextually within a conversation, as speakers accommodate each other. And, we can assume, both Didi and

Naomi are using “rich” in a way that is appropriate given their conversations. They are making no linguistic mistake, and we can assume that they know all of the relevant monetary and mathematical facts. Both, then, are speaking truly.

But the only way to reconcile (28-30) is to allow that the truth of a claim is relative to a context or conversation. For, if Didi disagrees with Naomi about whether Mary is rich, Didi must take Naomi’s claim (27) to be false. The very same claim, then, is true relative to Naomi’s context and false relative to Didi’s.

I am skipping over many interesting and subtle details, because what I want to discuss is not Richard’s relativism, but the limits he puts on its application. For, as we learn in Chapter 5, Richard thinks that it is a mistake to suppose that a relativist semantics can help make sense of disagreements of taste: cases where, although one party judges that a band is cool and the other that it is not cool, neither supposes that the other is mistaken. The basic argument is simple. If we take the claim that the band is cool to be a proposition with a relative truth value, then each party x will take the other to have asserted something false (relative to x ’s context). Since believing a falsehood is a way of making a mistake, neither can take the other party to be “faultless.” Richard concludes that, for a range of cases where we take there to be “faultless disagreement” about matters of taste, we should use a commitment semantics, not a relativist semantics.

This is puzzling, for a couple of reasons. First, “cool” is itself a gradable adjective, and so ought to exhibit the same kinds of contextual variation that “rich” and “flat” do. This is not to say that there are no differences, since there seems to be little contextual variation involving the comparative forms of “rich” and “flat,” whereas in the case of “cool” we have *both* variation in the ordering and variation in where the cutoff point lies. Still, even if “cool” exhibits some kinds of contextual variation not found in “rich,” it remains true that the kind of contextual variation Richard thinks relativism is needed to explain in the case of “rich” can *also* be found in the case of “cool.” Two people could agree about the ranking of bands from coolest to least cool, but disagree about whether a particular band is cool. Thus, if relativist semantics is really needed to understand this phenomenon in the case of “rich,” it ought to be needed too in the case of “cool,” whenever we want to say that two people disagree about whether something is cool, and regardless of whether the parties take each other to be mistaken.

Second, the reasons that might lead us to say, in the case of disagreements about what is “cool,” that the other party is not mistaken, are present as well in the disagreements about “rich” and “flat” that Richard discusses. Here is Richard explaining why he doesn’t think you’re mistaken in not thinking, as he does, that Sleater-Kinney is cool:

I think Sleater-Kinney is a terrific band, you find their music at best a cunning refinement on the sound of fingernails scraped upon

slate. When we explore our differences, I find that my reasons for liking them—the dissonance of minor keys married to Carrie Brownstein’s vocals speaks to me of barely suppressed rage—just don’t move you. You find the way the band conveys anger irritating. I can see why someone would feel that way. I don’t think you’re making a mistake in not being moved as I am; I acknowledge your attitude, though I don’t agree with it. (131)

But couldn’t I have the same attitude in a disagreement about whether someone is rich, or whether a field is flat? Richard seems to concede this, in a footnote:

Claims about wealth (and roundness, and height, and many other notions expressed by the gradable adjectives) may not always be bearers of truth. Suppose Chapter 4’s Didi looks upon her differences with Naomi about Mary as not involving any mistake on Naomi’s part. Suppose, that is, that Didi’s attitude is something along the lines of *That is a way of thinking about wealth; I don’t think it’s the best way to think about the matter, but it’s not as if it is flat-out mistaken*. Then the argument of the first few pages of this chapter may kick in: Didi disagrees with Naomi about *p*, but acknowledges that Naomi isn’t making a mistake. So Didi should not ascribe truth to *p*, or falsity to Naomi’s contrary opinion. (149 n. 35)

But then it’s hard to see how, given what Richard has said, the truth-relativist semantics could *ever* be appropriate. For the cases he describes in Chapter 4, in motivating the relativist approach, *are* cases where the parties to the debate are willing to say of each other: “That *is* a way of thinking about wealth; I don’t think it’s the best way to think of the matter, but it’s not as if it is flat-out mistaken.” Witness this passage, discussing a disagreement about whether Flanders Field is flat:

... you and I have common interests and aims—in particular, we want to practice soccer, to maximize the length of our practice, to avoid a long drive, to practice on something whose topography is close to that of a regulation soccer field. Our judgments about Flanders Field reflect a difference in the way in which we order the importance of these interests and aims. *Must* there be something about the situation, our interests and aims, and our judgments that would bring ‘any rational person’ after sufficient reflection and consideration to pronounce me correct about the field and you wrong, or vice versa? Surely not. We each want non-bumpiness and ease of access to the practice field; we are inclined towards different trade-offs in satisfying these desires. (113)

What is the difference, really, between Richard's attitude to your different priorities in this case, and his attitude to your different reactions, in the Sleater-Kinney case? If Richard's argument against a relativist treatment of "cool" is any good, it seems, then it should tell equally against a relativist treatment of "flat" and "rich."

It seems to me, though, that the argument has no force against a relativist treatment in any of these domains. Here's the argument, in a slightly more articulated form:

- (a) In some genuine disagreements over what is cool (or flat), the parties need not regard each other as making a mistake.
- (b) If the parties are disagreeing about a truth-evaluable proposition, each party x should take the other y to be saying something false (relative to x 's context).
- (c) If x should take y to be saying something that is false (relative to x 's context), then x should regard y as making a mistake.
- (d) So, in these cases, the parties are not disagreeing about a truth-evaluable proposition.

It seems to me that a relativist should simply reject (c). If it is plausible for a *nonrelativist* to call any assertion of a falsehood a mistake, presumably that's because of a bridge principle connecting truth to norms for assertion, something like

The Truth Norm: One should: assert that p only if p is true.

But presumably a Richard-style relativist will want a relativized version of this norm:⁸

The Truth Norm (Relativized): One should: assert that p only if p is true (relative to the context one occupies).

And this norm does not support (c); indeed, it is incompatible with it (on the assumption that some propositions are true only relatively).

Somewhat more concessively, the relativist might want to distinguish three senses of "making a mistake" in making an assertion:

⁸I have argued elsewhere (e.g., MacFarlane 2005) that if this is the *only* thing the relativist says connecting relativized truth to normative proprieties, then the relativist's position will be indistinguishable from a kind of "nonindexical contextualism," on which two parties can accept and reject the same proposition without disagreeing in any robust sense. To distinguish relativism from nonindexical contextualism, it is also important to say something about when assertions must be retracted. But since Richard leaves this issue unexplored, we need not get into it here.

- S is making a mistake₁ in asserting that p iff S 's assertion that p is unwarranted.
- S is making a mistake₂ in asserting that p iff S 's assertion that p violates the relativized truth norm. (Equivalently: iff p is not true relative to the context S occupies in making the assertion).
- S is making a mistake₃ in asserting that p iff it is false that p . (Here, "false" is used without explicit relativization; the proposition *that it is false that p* is true, relative to a context c , just in case p is false, relative to c .)

The relativist can now say that although the parties in our example should regard each other as making a mistake₃, they need not regard each other as making a mistake₁ or making a mistake₂. Premise (a) of Richard's argument is only plausible when "making a mistake" is interpreted in sense 1 or 2, while premise (c) is only compelling when "making a mistake" is interpreted in sense 3. Since there is no single interpretation of "making a mistake" that makes both (a) and (c) compelling, the relativist can reject the argument.

If these considerations are correct, then Richard is wrong to limit the applicability of relativist semantics in the way that he does; and the relativist technology can be applied in many areas where Richard thinks a commitment semantics is needed. Given the serious problems noted above with commitment semantics, I think that's a good thing.⁹

References

- Simon Blackburn 2001. *Ruling Passions*. Oxford: Oxford University Press.
- John MacFarlane 2005. "Making Sense of Relative Truth." *Proceedings of the Aristotelian Society* 105 (2005), 321–39.
- Mark Richard 2008. *When Truth Gives Out*. Oxford: Oxford University Press.

⁹This is a lightly revised version of my contribution to the Pacific APA Author Meets Critics session on *When Truth Gives Out*. I am grateful to Mark Richard for helpful correspondence.