This is a penultimate draft of a paper which is forthcoming in The Routledge Handbook of Philosophy of Relativism (ed. by Martin Kusch)

Assessment Relativism

Filippo Ferrari (Universität Bonn)

Abstract

Assessment relativism, as developed by John MacFarlane, is the view that the truth of our claims involving a variety of English expressions—'tasty', 'knows', 'tomorrow', 'might', and 'ought'—is relative not only to aspects of the context of their production but also to aspects of the context in which they are assessed. Assessment relativism is thus a form of truth relativism which is offered as a new way of understanding perspectival thought and talk. In this article, I present the main theses of assessment relativism, focusing in particular on highlighting the points of commonality and contrast with other forms of truth relativism. I then offer some critical remarks concerning the motivation of assessment relativism in relation to matters of taste.

1. Introduction

Assessment relativism (henceforth AR) is a type of truth relativism,[1] that has been developed by John MacFarlane in a series of works,[2] culminated in his 2014 book *Assessment Sensitivity: Relative Truth and its Applications*.

[1] See García-Carpintero and Kölbel 2008 for a variety of models of relative truth.

[2] The view is first sketched in MacFarlane 2003.

Relativism about truth is the thesis that (some) truths are true merely relatively. This view is mainly motivated by the attempt of making sense of the possibility of disputes where none of the competing opinions seems less legitimate, or *less true*, than the others. This phenomenon is known under the label "faultless disagreement" and, roughly put, it concerns situations where one party accepts while the other rejects that things are so-and-so but neither of them is, not even in principle, off-track and guilty of any mistake.[3]

To get a proper grip on AR and to distinguish it from other versions of truth relativism, three questions are particularly relevant: (i) which truths are relative and which aren't?—section 2; (ii) what is truth relativism and what are the bearers of (relative) truth?—section 3; (iii) in what sense is truth relative according to AR?—sections 4-6. Section 7 discusses two challenges to AR.

2. The scope of assessment relativism

AR is a *local* thesis targeting a restricted range of truths. Its exact range of application is ultimately an empirical question, having to do with whether and to what extent a given language—e.g. English—contains expressions which are apt for a relativistically-engineered notion of truth. MacFarlane takes 'tasty', 'knows', 'tomorrow', 'might', and 'ought' to be paradigmatic examples of expressions which are apt for AR. As a litmus test for detecting the presence of AR,

[3] Kölbel 2003 takes faultless disagreement as the main motivation for truth relativism. Rovane 2013 offers some criticisms to this strategy.

MacFarlane suggests to look not only at the phenomenon of faultless disagreement but also at that of retraction—a speech act targeting a previously made assertion which speakers perform by saying things like “I retract that”, “I no longer stand by that”. I will leave the question whether retraction has an adequate empirical support to experimental philosophers—see, for instance, Knobe and Yalcin 2014, Marques 2018, and Wright 2007—focusing on introducing, explaining, and assessing the theoretical framework of AR in relation to one specific application—matters of taste.

3. What is truth relativism?

To understand the mechanics of AR, we need first to understand the general thesis of relative truth. A natural starting point is to ask what are the bearers of (relative) truth. In this regard, a distinction between utterances, sentences, and propositions has to be made. Utterances are speech acts. They typically involve the use of a language and the intentional acts of speakers at a certain time and place. Sentences are linguistic expressions. They constitute the smallest unity of meaning that can be assessed for truth and falsity. Propositions are non-linguistic items that are taken to be the content of declarative uses of sentences. Clearly, two speakers can say the same thing by uttering different sentences, whether in the same or different languages. That *same thing* is the proposition expressed by those sentences. Sarah in uttering the English sentence “John is a philosopher” expresses the proposition that John is a philosopher; Elisabetta in uttering the Italian sentence “John è un filosofo” expresses the very same proposition expressed by the sentence uttered by Sarah.

So, what are the primary bearers of truth? It sounds infelicitous to say that a speech act of uttering—e.g. Sarah's *uttering* "John is a philosopher"—is true or false. Strictly speaking, an act of uttering might be correct or incorrect, appropriate or inappropriate, but it is not true or false.[4] Thus, we are left with either sentences or propositions to play the role of primary truth bearers. MacFarlane opts for propositions. One reason for this, as we will see, is that it's easier to capture the kind of relativity at the core of AR (MacFarlane 2014, ch.3).

Let's now introduce a standard piece of semantic machinery, the so-called "Kaplanian two-dimensional semantics". In his seminal paper "Demonstratives" (Kaplan 1989), operating within the framework of compositional semantics,[5] Kaplan distinguishes two aspects of meaning: character and content. The character of a linguistic expression—a sentence or a sub-sentential expression—reflects semantic rules governing how the content of that expression may vary from one context of use to the next. The sentence "I am a philosopher" contains the indexical expression 'I' which is governed by the following rule: 'I' refers to the user (speaker or writer). Such a sentence expresses different propositions in different contexts, depending on who's using it. As used by MacFarlane, it expresses the proposition that MacFarlane is a philosopher, while as used by Merkel it expresses the proposition that Merkel is a philosopher. These two

---

[4] For a dissenting voice, see García-Carpintero (this volume).

[5] Compositional semantics provides rules governing the interpretation of subsentential expressions and their modes of combination in order to explain how the meaning of whole sentences is determined by the meanings of their parts.

propositions might have different truth values: in fact, in the actual world, the former is true while the latter is false. Formally, the character of a linguistic expression is a function from contexts to contents—where a context, in its abstract sense, is a possible occasion of use of an expression that can be individuated as a sequence of parameters including an agent, a world, a time, a location. The second aspect of the meaning of a linguistic expression is the content. If the expression in question is a sentence, the content will be a proposition, whereas if it is a sub-sentential expression—e.g. a name or a predicate—the content will be an object, an individual, or a property. Formally, the content is represented by a function from circumstances of evaluation to an appropriate extension. By "circumstances of evaluation" Kaplan means

> [B]oth actual and counterfactual situations with respect to which it is appropriate to ask for the extensions of a given well-formed expression. A circumstance will usually include a possible state or history of the world, a time, and perhaps other features as well." (Kaplan 1989: 502).

The sentence "snow is white" expresses the same propositions in all contexts, but whether it is true will depend on the world at which it is evaluated: in the actual world the proposition expressed is true, but in a possible world where snow isn't white, it would be false.

What other information beside possible worlds should be included in the circumstances is "a matter of language engineering" (Kaplan 1989: 504). If we think, like temporalists do, that the sentence "Filippo is writing" expresses a temporally neutral content—i.e. a proposition which doesn't contain any information about the time at which Filippo is writing—then we might include a time-parameter in the circumstances in order to assess whether the proposition expressed by such a sentence is true or false. If we consider as the relevant value of the time parameter in the circumstances Wednesday, 26-12-2018 at 9am, then the proposition expressed is true, whereas if we consider Wednesday, 26-12-2018 at 3am, the proposition is false.

It is important to bear in mind that the use of a sentence always occurs in a context which contains information about the time and the world at which the sentence is used by an agent. However, if we believe that modally and temporally neutral contents should be allowed for, such information provided by the context does not leak in the proposition. Nevertheless, the context at which the sentence is used—let's call it, following Kaplan, *the context of use*—provides the default values of the relevant parameters (i.e. time and world in our example) in the circumstances of evaluation. Thus, if Filippo is uttering the sentence "Filippo is writing" in the actual world on December the 26th 2018 at 9am, the context in which such a sentence is used by Filippo will contain information about the world and the time at which it is used. If we then want to assess whether the proposition expressed by that sentence—i.e. the proposition that Filippo is writing—is true or false in relation to the context in which it is used, we fill the world and time values in the circumstances of evaluation with the information provided by the context of use. In this sense we can talk of *the circumstances of the context of use*.

The context of use thus plays two distinct roles, as MacFarlane puts it: a content-determining role—which proposition is expressed depends on features of the context of use—and a circumstances-determining role, selecting the circumstances of evaluation that are relevant to assess the truth of the proposition. To these two roles of the context of use correspond two ways in which a sentence can be context-sensitive: it can be use-indexical in that it expresses different propositions at different contexts of use (e.g. "I am a philosopher"); or, it can be use-sensitive in that it expresses the same proposition at all contexts of use but its truth depends on features of the circumstance of the context of use (e.g. "Filippo is writing" is, according to temporalists, use-sensitive but not use-indexical).

As we can now appreciate, that sentence truth is relative because the sentence is use-indexical, is a mundane kind of relativity. A more interesting relativity occurs once the proposition is fixed for all contexts, and yet its truth value might vary from one context to the next. If temporalism is true, my utterance of "Filippo is writing" as used by me on 26-12-2018 at 11am expresses the proposition that Filippo is writing which is true as evaluated at the circumstance of the context of use but false when evaluated at different circumstances of evaluation (e.g., on 26-12-2018 at 3am).

4. Non-indexical-contextualism as truth relativism

The idea of use-sensitivity without use-indexicality was adapted by some relativists to model the intuitive talk of relative truth in some domains—e.g., that of taste (see Egan 2007, Kölbel 2002, Recanati 2008, Richard 2008). The thought is to take Kaplan's framework and to enrich the circumstances of evaluation with a taste parameter tracking the gustatory sensibility of the speaker. According to this model—known as Non-Indexical-Contextualism (NIC for short)—an utterance of a sentence containing a taste predicate—e.g. 'oysters are tasty'—expresses the same proposition in all contexts of use—i.e., the proposition that oysters are tasty—but its truth-value varies from one context of use to the next in tandem with variation in the agent's gustatory sensibility. Thus "oysters are tasty" would invariably express the proposition that oysters are tasty which is true as used by me but false as used by my brother. In short, according to NIC the truth of taste propositions is relative to the gustatory sensibility operating at the context of use. In this way NIC effectively accounts for faultless disagreement. If, in claiming "these oysters are tasty", I endorse the proposition that these oysters are tasty while my brother, in claiming "nah, these oysters are not tasty", endorses the proposition that these oysters aren't tasty, we would disagree but each of us would say something true relative to our respective contexts of use. Because there's nothing wrong in endorsing a true proposition, there's a clear sense in which our disagreement is faultless.

Whether the truth relativity exemplified by NIC deserves the honorific title of "genuine truth relativism" is open to discussion. MacFarlane thinks that since NIC preserves an element of absolutism it doesn't count as a thoroughgoing form of relativism. Once we take the proposition used by an agent and we pair it with the context of use we get a once-and-for-all truth value. This is sometimes expressed, perhaps misleadingly, by saying that within NIC utterance truth is absolute

(MacFarlane 2014: 107).[6] Take the sentence "oysters are tasty" as uttered by me. Given that there is one privileged context of use—that of my utterance—there's a privileged default value for the taste parameters in the circumstances of evaluation. In this sense, the question whether the proposition expressed by the sentence "oysters are tasty" as used by me is true has a unique and absolute answer.

5. Assessment Relativism

If we want to cross "the philosophically interesting line between truth absolutism and truth relativism [we need to] relativize truth not just to a context of use [...] but also to a context of assessment" (MacFarlane 2014: 60). Thus, what makes a view about truth a genuinely relativistic one is not the mere addition of a special parameter in the circumstances of evaluation—e.g., a taste parameter—but rather the addition of a new context besides the context of use—i.e., a context of assessment. Such a context, which is structurally analogous to the context of use,

[6] MacFarlane can express this point without using the notion of utterance truth but that of accuracy instead: "an attitude or speech act occurring at c_1 is accurate, as assessed from a context c_2, just in case its content is true as used at c_1 and assessed from c_2" (MacFarlane 2014: 127). For NIC an utterance made by Ben of "oysters are tasty" in context c_1 and an utterance of "oysters aren't tasty" made by Bob in c_2 can both be accurate as assessed from any context. AR predicts that the accuracy of Ben's utterance precludes the accuracy of Bob's utterance. Thus, utterance's accuracy is invariant in NIC but not in AR.

is a possible situation in which a use of a sentence might be assessed. Its primary function is to provide the value for the special parameters.

This gives the semantic machinery an interesting twist. In fact, AR can be defined as the thesis that the truth of certain propositions—i.e. those that are expressed by a use of sentences containing assessment-sensitive expressions—is relative not only to aspects of the context of use but also to aspects of a context of assessment. Focusing on taste propositions, MacFarlane characterizes assessment-relative truth as follows:

> REL — A proposition p is true as used at the context of use (c_1) and assessed from a context of assessment (c_2) if and only if p is true at $<wc_1,gc_2>$, where wc_1 is the world of c_1 and gc_2 is the taste of the agent of c_2 (the assessor). (MacFarlane 2014: 105)

While the job of c_1 is that of determining the proposition expressed, that of c_2 is to provide the values for the gustatory parameter (g) in the circumstances of evaluation. To illustrate: I'm at the restaurant with my brother and I claim: "oysters are tasty". Here the context of use (c_1) is the context in which I utter the sentence "oysters are tasty". It determines the proposition expressed—namely the proposition that oysters are tasty—and it fixes the default value of the world parameter in the circumstances of evaluation (wc_1)—i.e., the actual world. Since there is no use-indexical expression, such a content remains invariant across contexts. The role of a context of assessment (c_2) is that of providing the value for the taste parameter in the circumstances of evaluation (gc_2) which determines whether the extension of 'tasty' includes the oysters (making the proposition true) or not (making it false). Since I like the oysters while my brother dislikes them, the proposition that oysters are tasty is true as assessed from my context of assessment but false as assessed from my brother's context of assessment. In this way, AR can account for the possibility of faultless disagreement.[7]

REL shows the formal difference between AR and NIC. While the two views agree that an utterance of "oysters are tasty" invariantly expresses a single proposition, they disagree about what gustatory standard is relevant for assessing the truth of the proposition expressed. According to NIC the default value for the relevant standard is provided by the context of use whereas according to AR it is provided by a context of assessment. At first sight, this might look an insignificant difference. On a closer look, though, it has important consequences since it frees AR from the residual element of absolutism that affects NIC. Given that there's no privileged context of assessment which sets the default value of the taste parameter in the circumstances of evaluation, in AR utterance truth isn't absolute. Even keeping fixed the context of use and the proposition expressed, we don't get a once and for all truth value.[8] Moreover, relativising truth to both the context of use and a context of assessment has important normative consequences, to which we now turn.

6. The normative profile of assessment relativism

[7] While AR and NIC agree on the relevant notion of faultlessness, they work with different notions of disagreement. See MacFarlane 2014: 132.

[8] See footnote 6 for an elaboration of this point.

Dummett argued that to have a proper grasp of what truth is, we need not merely to know under what circumstances propositions are true, but also to understand the connection between their truth and the proprieties of their use (Dummett 1959). Following Dummett, MacFarlane discusses at length the practical difference that his theory makes with respect to the making and retracting of assertion. Starting with the idea that truth is the core—not necessarily the sole—normative notion of assertion, MacFarlane subscribes to the following reflexive truth norm:

> RTN — An agent is permitted to assert that p at context c_1 only if p is true as used at c_1 and assessed from c_1. (MacFarlane 2014: 103)

RTN is a bridge principle linking semantics and pragmatics. It says that the truth of an assertion made at c_1 and assessed from c_1 is a necessary condition for the permissibility of the asserting at c_1. However, if RTN were the only normative principle there would be no practical difference between AR and NIC since the assessment context plays no distinctive role there: whenever we make an assertion in a context we are also assessing it from that same context. NIC and relativism turn out to be normatively equivalent theories according to RTN. This means that RTN doesn't give us the full story concerning the normative significance of AR.

It is only when we turn to a different conversational phenomenon—that of retracting—that the operational difference between NIC and AR is appreciable. By 'retraction' MacFarlane means the speech act one performs in saying 'I take that back' or 'I retract that'. The target of retraction is a previously made but unretracted speech act (e.g., an assertion). As any other speech act, retraction has its distinctive normative profile which MacFarlane characterizes as follows (Retraction Norm):

> RN An agent in context c_2 is required to retract an (unretracted) assertion of p made at c_1 if p is not true as used at c_1 and assessed from c_2. (MacFarlane 2014: 108)

RN is a prescriptive norm—it says that the untruth of a proposition as used at c_1 and assessed from c_2 is sufficient to require retracting a previous assertion of it. The effect of retracting is that of rendering the previous assertion null and void—to disavow the assertoric commitments undertaken in the original assertion.

With both RN and RTN on board, we can fully appreciate the practical difference between NIC and AR. Because NIC does not distinguish operationally between a context of assertion and a context of assessment, when it is equipped with a retraction norm it predicts that a subject ought to retract a previously unretracted assertion just in case that assertion was impermissible—i.e. it expressed a false proposition in the context in which it was performed. By contrast, AR predicts that a subject is required to retract an assertion whenever the proposition it expresses is false from her current context of assessment regardless of whether such an assertion was deemed permissible by RTN. To illustrate: suppose that in c_1 I like oysters and I assert "oysters are tasty" but in c_2 I change my tastes and I no longer enjoy oysters. AR—but not NIC—predicts that at c_2 I am required to retract my assertion of "oysters are tasty" made at c_1 since the proposition it expressed is assessed as false in c_2, even though that assertion was permissible in c_1 according to RTN.

7. Assessing assessment relativism

AR is an ingenious proposal which sets a milestone in the history of relativism. The view is, however, not immune to criticisms. Among the many objections that have been raised against AR,[9] I'll just mention two issues in relation to retraction and faultless disagreement.

Recall, an act of retraction targets a previously made speech act—e.g. an act of asserting—not its content. However, according to RN, what triggers the requirement to retract are characteristics of the asserted proposition, namely its untruth relative to the context of assessment where the retraction takes place. By retracting, an agent undoes the normative changes effected by her assertion in the original context with the result that she "is no longer obliged to respond to challenges to the assertion […], and that others are no longer entitled to rely on [her] authority for the accuracy of this assertion".[10] However, in retracting an agent is not required to admit fault, and for two reasons: first, given that RN states that the untruth of a proposition is only a sufficient condition for the obligatoriness

[9] Just to mention a few criticisms: for objections against the application of AR to (i) epistemic modals see Ross and Schroeder 2013, Von Fintel and Gillies 2008, Wright 2007; (ii) to future contingent see Moruzzi and Wright 2008, Brogaard 2008, Dietz and Murzi 2013; (iii) to knowledge ascription see Brendel 2014, Kompa 2002, and Wright 2017.

[10] MacFarlane 2014: 108.

of retracting, nothing precludes an agent to retract an assertion which is true as assessed from her context of assessment. Second, even when a subject is required to retract, the retracted assertion might have complied with RTN when it was made in the original context of assertion. In both cases, an attribution of fault is inappropriate.

Do we need AR and the specific sense of retraction that MacFarlane has in mind to obtain this normative effect? I'm inclined to agree with Marques 2018 and Raffman 2016 that we do not. Let's consider an example concerning matters of taste: until the age of twenty I didn't like the taste of mustard and, as I recall now, I've asserted several times "mustard is disgusting". Then my taste sensibility changed and I started liking mustard. During a dinner at which my mother has prepared mustard, noticing, with surprise, that I'm the first to grab some mustard, she claims: "Didn't you think mustard is disgusting?". To that I reply: "Well, I've changed my mind, and I no longer stand by my previous assertion". Let's label this speech act *withdrawing*. Everybody at dinner would take my withdrawing to signal my intention of distancing myself from my previous assertions that mustard is disgusting and in fact to withholding from the assertoric commitments engendered by my previous assertions. This, as expected, does not involve an admission of fault for having asserted in the past that mustard is disgusting: since I disliked mustard, I consider my previous assertions as perfectly reasonable.

What I did with my withdrawing is thus: (i) to deny that mustard is disgusting, and (ii) to manifest my disposition to withhold from the assertoric commitments associated with my previous assertions. It thus seems that my withdrawing would achieve the normative revisions that characterize an act of retraction without, however, requiring AR. An advocate of NIC, for instance, could easily account for the normative effect of my withdrawing. The challenge for an advocate of AR is thus to individuate some normative effects distinctive of retraction that are not already carried out by an act of withdrawing.[11]

The second issue concerns what I take to be an ingredient, and important, part of the folk conception of faultless disagreement, namely the idea that the faultlessness of certain kinds of disagreement can be appreciated and coherently expressed not just from the abstract point of view of the formal framework of AR, but also from within a committed perspective taking part to the dispute. Wright calls this extra ingredient *parity*, and according to him it is "meant to be implicated by faultlessness—conveyed in the acknowledgment that your opinion is just as good as mine".[12] For this reason, intuitions about parity are arguably explanatorily prior to intuitions about faultlessness: it is because disputants think that their opinions are roughly on a par that they think that the disagreement between them is faultless. A fully satisfactory account of the phenomenon of faultless disagreement must include an account of the parity ingredient.

Can AR account for it? In its current form, it does not seem that it can. The gist of the argument is the following. Suppose I disagree with John about whether oysters are tasty: I claim they are tasty while John claims they are not. My assessment context is one in which the proposition that oysters are tasty is true: relatively to

[11] See Ferrari 2016b for a criticism of MacFarlane's take on retraction.

[12] Wright 2012: 439.

that context the proposition that oysters aren't tasty is assessed as false. But that's exactly what John has asserted. So how can I judge from within my context of assessment that John's assertion is on a par with mine if what I've claimed is true while what John has claimed is false—as assessed from my context of assessment?[13] If an attribution of falsity carries some normative weight—as it should do[14]—it seems that I cannot but assess John's claim to be *not* on a par with mine: after all it expresses a proposition I'm committed to assess as false. It is thus clear that the framework of AR, as it stands, does not allow for the possibility of expressing equal legitimacy of contrary opinions within a committed perspective. This isn't, of course, meant to be a conclusive objection to AR: it just illustrates that more work has to be done by an advocate of AR both on the normative functions of truth and falsity and on the metaphysics of taste.[15, 16]

---

[13] Cf. MacFarlane 2014, pp. 106-107.

[14] Cf. Wright 1992, Ch. 1.

[15] Ferrari 2016c develops some suggestions for how to improve AR with a pluralist conception of the normative function(s) of truth. Ferrari 2016a and Ferrari 2018 gove a full account of normative alethic pluralism in relation to a comparative analysis of disagreement in different domains.

[16] I'd like to thank Elke Brendel, Manuel García-Carpintero, Andrea Crepoli, Martin Kusch, Teresa Marques, Sebastiano Moruzzi, Stefano Pugnaghi, Elisabetta Sassarini, and Dan Zeman for their useful feedback. Moreover, I'd like to acknowledge the support of the Deutsche Forschungsgemeinschaft (DFG—BR 1978/3–1) for sponsoring my postdoc at the University of Bonn. While working on this paper, I have benefitted from participation in the Pluralisms Global

References

Brendel, Elke (2014) “Contextualism, relativism, and the semantics of knowledge ascriptions”, *Philosophical Studies* 168:101–117.

Brogaard, Berit (2008) “Sea Battle Semantics”, *The Philosophical Quarterly* 58(231): 326-335.

Dietz, Richard and Murzi, Julien (2013) “Coming True”, *Philosophical Studies* 163:403–427.

Dummett, Michael (1959) “Truth”, *Proceedings of the Aristotelian Society*, n.s. 59: 141–62.

Egan, Andy (2007). “Epistemic Modals, Relativism and Assertion”, *Philosophical Studies* 133: 1–22.

Ferrari, Filippo (2016a) “Disagreement about taste and alethic subrogation”, *Philosophical Quarterly* 66(264): 516-535.

Ferrari, Filippo (2016b) “Assessment Sensitivity”, Analysis 76(4): 516–527.

Ferrari, Filippo (2016c) “Relativism, Faultlessness and Parity”, *Argumenta*, 3: 77-94.

Ferrari, Filippo (2018) “Normative Alethic Pluralism” in J. Wyatt et al. (eds.), *Pluralisms in Truth and Logic*, London & New York: Palgrave MacMillan: 145-168.

García-Carpintero, Manuel (2019) "Semantics and Metaphysics of Truth", this volume.

Research Network (National Research Foundation of Korea grant no. 2013S1A2A2035514). This support is gratefully acknowledged.

García-Carpintero, Manuel and Kölbel, Max. (2008) (eds.) *Relative truth*, Oxford: Oxford University Press.

Kölbel, Max (2002) *Truth Without Objectivity*, London: Routledge.

Kölbel, Max (2003) "Faultless Disagreement", *Proceedings of the Aristotelian Society*, 105: 53–73.

Kompa, Nikola (2002) "The context sensitivity of knowledge ascriptions", *Grazer Philosophische Studien*, 64: 1-18.

MacFarlane, John (2003) "Future contingents and relative truth", *Philosophical Quarterly*, 53(212): 321–336.

MacFarlane, John (2014) *Assessment Sensitivity: Relative Truth and its Applications*, Oxford: Oxford University Press.

Marques, Teresa (2018) "Retractions", *Synthese* 195:3335–3359.

Raffman, Diana (2016) "Relativism, Retraction, and Evidence", *Philosophy and Phenomenological Research* 92(1): 171-178.

Ross, Jacob and Schroeder Mark (2013) "Reversibility or Disagreement", *Mind*, 122: 43-84.

Recanati, François (2008) "Moderate Relativism", in García-Carpintero and Kölbel: 41-62.

Richard, Mark (2008) *When Truth Gives Out*, Oxford: Oxford University Press.

Rovane, Carol (2013) *The Metaphysics and Ethics of Relativism*, Cambridge(MA): Harvard University Press.

Von Fintel, Kai and Gillies, Anthony (2008) "CIA Leaks", *Philosophical Review*, 117(1): 77-98.

Wright, Crispin (1992) *Truth and Objectivity*, Cambridge(MA): Harvard University Press.

Wright, Crispin (2007) “New Age Relativism and Epistemic Possibility: The Question of Evidence”, *Philosophical Issues* 17(1):262-83.

Wright, Crispin (2012) "Replies part III: truth, objectivity, realism and relativism", in Coliva (ed.) Mind, Meaning and Knowledge: Themes from the Philosophy of Crispin Wright, Oxford: Oxford University Press: 418-450.

Wright, Crispin (2017) “The variability of ‘knows’: An Opinionated Overview”, in Jonathan Ichikawa (ed.), *Routledge Handbook of Epistemic Contextualism*: 13-31.