The Creative Interpreter: Content Relativism and Assertion¹

Herman Cappelen

Draft, September 11, 2008

Philosophers of language and linguists tend to think of the interpreter as an essentially *non-creative* participant in the communicative process. There's no room, in traditional theories, for the view that correctness of interpretation depends in some essential way on the interpreter. As a result, there's no room for the possibility that while P is the correct interpretation of an utterance, u, for one interpreter, P* is the correct interpretation of that utterance for another interpreter.

Recently, a number of theorists have, for separate reasons, argued in favour of a radically different view of communication – a view in which the interpreter and her context play what should be thought of as *a content-creating role*. According to such views, natural languages contain what I'll call *interpretation sensitive terms*: terms the correct interpretation of which varies across interpreters (or, more generally, contexts of interpretation).³ An interpretation sensitive sentence can have one content relative to one interpreter and another content relative to another interpreter. This paper is a development and (partial) defence of the view that interpretation sensitivity is ubiquitous in natural language. I call the view that there are interpretation sensitive terms *content relativism*.

Before starting the discussion of content relativism, it is worth pointing out that recent attempts to develop semantically motivated versions of truth relativism should be seen as part of this trend of giving the interpreter a more active role.

¹ This paper started out as part of joint project with Crispin Wright (triggered by our work in Arché's Contextualism and Relativism Project) and his input along the way has been invaluable. I'm grateful for comments from and discussions with Jessica Brown, Timothy Endicott, Torfinn Huvenes, Ernie Lepore, Jed Lewinsohn, Stephen Neale, François Recanati, Jonathan Schaffer, Anders Schoubye, Andreas Stokke, Rachel Sterken, Brian Weatherson, and Elia Zardini,. The paper was much improved by Federico Luzzi's editorial assistance. Versions of the paper has been presented at several workshops and conferences, including the World Congress of Philosophy in Korea, Semantics and Philosophy in Europe at IHPST, Paris, Language and Law at CSMN/UiO, Arché's Contextualism and Relativism workshop, Glasgow University, and the Logos Research Group in Barcelona. I'm grateful to audiences on those occasions for helpful questions and discussions. The paper builds on work done in Cappelen (2008) and develops ideas found in various papers and books by Cappelen and Lepore (e.g., (1997) and (2004)).

³ Contexts of interpretation are just what you'd think: contexts in which utterances are interpreted. I won't choose here between different accounts of what contexts are - n-tuples of parameters, concrete situations, centred worlds, or what have you. For the most part, you can operate throughout with your favourite account of what a context is (if you have one) and just replace the speaker with an interpreter.

According to truth relativists⁴, the interpreter contributes not a component of content, but, rather, a parameter of the circumstance of evaluation. Here's a simple illustration of the difference between truth and content relativism, applied to claims involving the evaluative term 'yummy'.⁵

- According to truth relativists, claims about yummyness are evaluated for truth or falsity relative to standards of taste (or yummyness).⁶ They can't be true or false *simpliciter*. The proposition *that A is yummy* can be true relative to one standard of yummyness and false relative to another. According to truth relativists, the content expressed by utterances of sentences containing 'yummy' makes no reference to a standard of yummyness. Such contents are standard-neutral.
- A content relativist, on the other hand, thinks that standards of yummyness do enter into the content expressed by utterances of sentences containing 'yummy'. In contrast to traditional contextualists, they don't think these standards are fixed in the context of utterance. They are, instead, fixed in the context of interpretation. Relative to one interpreter, an utterance of 'A is yummy' can express the proposition *that A is yummy by standard S*, and relative to another interpreter it can express the proposition *that A is yummy relative to standard S** (where S and S* are the standards fixed in the interpreters' respective contexts).

Both views diverge from traditional theories by seeing the interpreter as contributing an important component to the communicative process. They differ in what that component is. For reasons spelled out in Cappelen and Hawthorne's *Relativism and Monadic Truth*, I think truth relativism fails – I won't rehearse those arguments here. The focus, instead, is entirely on content relativism. This paper has four sections:

- (i) In Section One, I show how content relativism can be embedded in a slightly revised version of David Kaplan's framework described in 'Demonstratives'.
- (ii) In Section Two, I briefly present four examples of how content relativism can be put to work. The first three draw on some recent work by Andy Egan (on indexicals), Brian Weatherson (on conditionals), and von Fintel & Gillies (on epistemic modals). My primary focus is on a content relativistic account of instructions (and, more generally, in imperatives and legal texts) such as 'Move suspicious packages away from crowds'. Neither case is explored in great detail since the goal of this section is to prepare for sections Three and Four.
- (iii) Section Three is in part taxonomical: it is unclear from the recent literature (and from the illustrations given in Section Two) what the relationship is between content relativism and three other views (or phenomena): speech act pluralism, utterance pluralism, and contextual vagueness. The goal of this section is to show ways in which the views are closely related and to make some suggestions for how they can be distinguished.
- (iv) In Section Four, I address what I take to be the main source of resistance to content relativism: that it allegedly can't make sense of assertion.

⁴ See for example MacFarlane (2005) and Egan, Hawthorne, Weatherson (2005).

⁵ 'Yummy' is appealed to only for illustrative purposes, I don't think either truth or content relativism are plausible account of 'yummy', see e.g. chapter four of Cappelen and Hawthorne (forthcoming).
⁶ See for example Lasersohn (2005), Egan, Hawthorne, and Weatherson (2005), and MacFarlane (2005).

Responding to this challenge is particularly important since, if successful, it would effectively block all attempts to implement content relativism.

1. Content Relativism as Interpretation Sensitive Characters

Kaplan describes characters as functions from contexts of utterance to contents.⁷ Context sensitive terms, e.g., 'I', have variable functions – i.e., the content varies with the context of utterance. Context insensitive terms, e.g., proper names, have stable characters, i.e., the function from a context of utterance to the semantic value always yields the same object.

One natural way to implement interpretation sensitivity, the one primarily explored in this paper, is to construe (at least some) characters as functions from pairs consisting of a context of utterance (Cu) and a context of interpretation (Ci) to contents. This leaves us with four possible categories of terms:

- (1) Terms that are stable with respect to all (Cu, Ci) pairs, e.g., proper names.
- (2) Terms that vary only with respect to the Cu component and have no variability when Ci is changed. The expressions discussed in David Kaplan's 'Demonstratives' are often taken to belong to this class.⁸
- (3) Terms that vary with respect to both Cu and Ci.
- (4) Terms that vary only with respect to Ci and not with Cu.

The focus in this paper is on category (3) terms. While there might well be category (4) terms, they will not be discussed in this paper.⁹

Four general points are worth emphasizing about interpretation sensitivity before moving on, in Section Two, to illustrations of various ways in which a content relativistic framework can be implemented.

(i) Explicit and Unarticulated Interpretation Sensitivity

Since it is expressions that have characters, my proposed construal of interpretation sensitivity applies only to lexical items. That might not seem like much of a restriction, but it rules out a certain kind of interpretation sensitivity that I don't mean to exclude. The basic idea I want to explore doesn't depend on there being a syntactic trigger. Two kinds of metaphors are often invoked to present an alternative picture. Some philosophers like to talk about 'gappy' propositions – they're like propositions, but have little holes in them. If you put the right kind of thing into the hole, you have a complete proposition. Others prefer the metaphor of a propositional 'skeleton' with missing limbs that, when added, give you a complete skeleton, i.e., a complete proposition. No matter which metaphor is invoked, the underlying thought is the same: the content expressed is not fully propositional. The speaker intends for the

⁷ See Kaplan (1989). Kaplan sometimes talks of contents, what he equates with what-was-said, as intensions, i.e., functions from circumstances of evaluation to truth values. Sometimes he talks about them as structured propositions. For the purposes of this paper, there's no need to choose a particular theory of content.

⁸ See Andy Egan's discussion of 'you' for some reservations about this (Egan (forthcoming)).

⁹ I think many terms traditionally classified as category (1) terms are really category (4) terms. That won't be argued for in this paper.

content to be completed; some kind of pragmatic process takes us from the incomplete content to a fully propositional content.¹⁰

Interpretation sensitivity can be combined with the view that (some) expressed contents are skeletons. On this view, the missing limb is replaced, not in the context of utterance, but in the context of interpretation, and the correct replacement can vary across contexts of interpretation (in gap-speak: the gap is filled in contexts of interpretation, and the correct filling varies across contexts of interpretation). In what follows I focus on cases where the interpretation sensitivity can be traced to a syntactic element, but I don't mean to rule out the possibility of a skeleton version of the view, and it should be quite obvious how to extend the view to such cases.¹¹

(ii) Range of Admissible Contexts of Interpretation

I will assume throughout that interpretation sensitive characters will appeal to the notion of a **range of admissible contexts of interpretation**. The basic idea is simple. Let u be an utterance of an interpretation sensitive term, I_i. There will be a range of admissible contexts of interpretation $<Ci_1...Ci_n>$ in which we can fix the content of I_i. Proponents of interpretation sensitivity *need* not hold that *all* contexts of interpretation are permissible. Of course, the range *could* be unrestricted, but it need not be. What determines the range of permissible Ci's? I'll assume that this is fixed in the context of utterance by some complicated mechanisms that centrally involve the speaker's intentions (and where these intentions interact in complicated and poorly understood ways with other aspects of the context of utterance).¹² It would be good to have something more specific to say about these mechanisms, but I don't.

A potential interpreter of an interpretation sensitive sentence has to figure out whether she's an admissible interpreter. This makes salient a question that has no correlate in interpretation *in*sensitive languages: what do we say about those interpreters who find themselves in contexts of interpretation outside the admissible range? I suspect there might be no unified answer to this question and that the status of inadmissible interpreters will vary across contexts and the kinds of expressions involved. One generalization might be true: the interpretations provided by interpreters in permissible contexts often have a different normative status from those offered by non-admissible interpreters. Their interpretations often override, in ways to made precise, those of non-authorized interpreter. I give some illustrations of how this can be cashed out in the case of instructions and laws (and other legal language) in Section Three below.

(iii) Mixed Interpretations

One feature of interpretation sensitive sentences that will be of some importance later is their potential for what I'll call 'mixed interpretations'. Suppose sentence S1 is interpretation sensitive and that S2 has no interpretation sensitive components, but

¹⁰ For one version of this view, see Kent Bach (1994).

¹¹ That said, I should point out that I don't think that these metaphors of gaps and missing limbs can be spelled out in ways that will ultimately prove to be useful in a theory of communication. I include them here simply because I think it is important to make clear that content relativism is a view that could be easily endorsed by those who embrace skeletons.

¹² I should mention another, more radical view according to which the range of admissible Ci's varies across Ci's. This rather dizzying view will not discussed in this paper, but I do consider it an important option when more detailed views of interpretation sensitive views are developed.

does have components that are sensitive to the context of utterance (e.g., contains an ordinary indexical such as 'I'). The content expressed by an utterance of the conjunction of S1 and S2 is fixed in part in the context of utterance and in part in the context of interpretation. The interpreter has to use some features of the context of utterance and some of the context of interpretation to fix the content expressed by a single utterance. This will be important when discussing utterance pluralism in Section Three below.

(iv) Interpretation Sensitivity: Semantic or Pragmatic?

I've talked about interpretation sensitive *contents* without specifying whether these contents are semantic, pragmatic, or some other kind of creature. This is in part because I am suspicious of the usefulness of any kind of broad semantics/pragmatics distinction (since it covers up too many useful distinctions¹³), and also because the general mechanism I am interested in can be described without commitment to any particular view of how such a distinction is best drawn. What matters for the arguments in this paper is only that there's an important level of asserted content that is best described as interpretation sensitive. Whether you choose to use the word 'semantic' to describe that level doesn't much matter in what follows.

For those who like to frame their thinking about these issues in terms of a semantics/pragmatics distinction, it is natural to construe interpretation sensitivity as I have described it above as a semantic mechanism. On most views, Kaplan's characters fix semantic content, and I've presented interpretation sensitivity in Kaplanian terms. It would, however, be easy to include interpretation sensitivity in an account of some non-semantic level of communicated content. This is particularly easy to do for those philosophers who sharply distinguish semantic content from speech act content.¹⁴ Such philosophers construe what is said (and asserted) by an utterance as being on the non-semantic side of communicated content and are therefore free to treat semantic content as interpretation *in*sensitive even though asserted content varies from one interpreter to another.¹⁵

2. Illustrations of Content Relativism at Work

In this section I first give three brief illustrations of how content relativism has been put to use: a) as an account of epistemic modals (Fintel and Gillies (2008)); b) as an account of indicative conditionals (Weatherson (forthcoming)); and c) as an account of the use of certain indexicals (Egan (forthcoming)). The goal of these illustrations is just to give the reader a sense of how content relativism can and has been put to use, not to evaluate or present the arguments given for each of these views (though it's interesting to note that these arguments are all of very different kinds—they're not based on one 'master argument' the refutation of which would suffice to undermine the view.) After these three brief illustrations, I turn to one case not yet discussed in the literature: I present a slightly more elaborate case for content relativism as a model of how we intend that certain instructions, imperatives, and legal texts should be interpreted.

These are meant as illustrations of how content relativism can and has been

¹³ See Cappelen (2007).

¹⁴ See e.g. Cappelen and Lepore (1997) and (2004), Soames (2002), and Salmon (1991).

¹⁵ For an extended defense of this kind of view, see Cappelen and Lepore (2004): Chapters 10-13.

put to use. I don't dwell extensively on any one of them - it goes beyond the scope of this paper to establish that content relativism provides the best explanatory framework in all these domains. The primary goal of this paper is to make clear what content relativism is, explain its complicated relationship it to some closely related views and, in Section Four, to respond to what I take to be the central source of opposition to it as a legitimate explanatory model: that it is incompatible with our best theories of assertion.

Illustration 1: Weatherson on Indicative Conditionals

According to Weatherson (forthcoming) the content of 'If p, q' is of the form: C(p, q, X). Weatherson treats 'X' as a plural variable and its values are '[...]the 'background' propositions relative to which the conditional is assessed'. C is the conditional relation and '[...]the value of X is (by default) R(x), where r is some epistemic relation (on a broad construal of 'epistemic') and x is a salient individual.' (12)

Weatherson's view is a version of what I call 'content relativism' because the salient individual (the one used to fix background propositions that constitute the value X in 'C(p, q, X)') is the interpreter (Weatherson talks about 'assessor'), not the speaker. Weatherson says:

The indexical relativist position [this is Weatherson's term for his what I call 'content relativism'] is that the content of an utterance of a conditional is (by default) a propositional frame that we might express as C(p, q, R(PROJ)). Relative to an assessor *a*, the content is C(p, q, R(a)).¹⁷ (*ibid*.)

The value of X will vary with the interpreter, and so the proposition asserted by an utterance of 'If P, Q' can vary between interprets.

Weatherson's argument for this view has two components. He first argues that an old-fashioned contextualist has no good explanation of why it is so easy to agree with others utterances of conditionals. Second, he argues that a content relativistic account of the indicative conditional is better placed to explain what goes on in McGee's alleged counterexamples to *Modus Ponens*. It goes beyond the scope of this paper to evaluate Weatherson's arguments, but even if it turns out that this isn't the final word on indicative conditionals, Weatherson has at least shown that content relativism is an explanatory framework that should be considered in any seriously discussion of the topic.

Illustration 2: Fintel and Gillies on Epistemic Modals

Kai von Fintel and Anthony Gillies (2008) are primarily concerned with refuting truth relativistic accounts of epistemic modals (e.g. the view defended by MacFarlane, Egan, Hawthorne, and Weatherson). At the end of their paper they briefly present their own view and, at least on one reading, it is a version of content relativism.

 ¹⁷ Weatherson's treatement of 'PROJ' is in part inspired by Tamina Stephenson's treatment of epistemic modals in Stephenson (2007).
 ²⁰ For a classic discussion of these issues in philosophy of law, see e.g. H.L.A. Hart's *The Concept of*

²⁰ For a classic discussion of these issues in philosophy of law, see e.g. H.L.A. Hart's *The Concept of Law* (in particular Chapter Seven). For a more recent treatment or related issues, see T. Endicott *Vagueness and Law*.

According to Fintel and Gillies, utterances of sentences containing epistemic modals are open to a variety of admissible interpretations. Someone who uses an epistemic modal has chosen to not explicitly restrict the claim to the evidence available to her (or someone else) at the time of utterance. In so doing, she has allowed that '[...] her sentence could be interpreted in a variety of ways: as a solipsistic claim about her current evidence, as a claim about what a group to which she belongs currently knows, or even as a claim about all the evidence available to (but not necessarily already processed by) that group.' (96) These are all, according to Fintel and Gillies, in some sense correct interpretations of her utterances. They appeal to the metaphor of a 'cloud of admissible contexts' – each context in the cloud yields a different interpretation. In so far as each context yields an admissible interpretation and the speaker intends for all of them to be admissible interpretations, this is an instance of what I call content relativism.

I should mention that this isn't the only (and maybe not the most natural) interpretation of what von Fintel and Gillies have in mind, and they don't explicitly present their view as a version of content relativism. Two interpretative complications are worth mentioning: First, in the last pages of their paper, they describe their view as one that appeals to 'contextual vagueness'. In Section Three below, I discuss the relationship between content relativism and contextual vagueness and I argue that some instances of the latter (including von Fintel and Gillies') should be classified as instances of the former. Second, they endorse a norm of assertion for epistemic modals according to which a speaker only needs to be justified in asserting the proposition that limits the evidential base to the speaker's 'own state of evidence' (21). They say, 'Of course, the hearer might interpret the sentence in a stronger fashion-and the speaker may well have anticipated and intended that-but the speaker can always retreat to the weakest interpretation, so that interpretation is the one she needs to be able to defend." (ibid.) Here is a problem for this view: Suppose the 'cloud of admissible contexts' consists of three contexts, C1-C3 and that C1 is the 'weakest', i.e., the one where the epistemic base is limited to the speaker's own state of evidence. Suppose also that the speaker chooses to retreat to C1. In what sense, then, are C2 and C3 admissible? In presenting their view as an instance of content relativism I am taking seriously the idea that these other contexts are admissible and I take that to imply that the speaker has asserted the contents expressed relative to these other contexts. These issues will be pursued further in Section Four below.

Illustration 3: Egan on 'You' and Other Indexicals

In a forthcoming paper, Andy Egan considers the view that indexicals should be given a content relativistic treatment. One example he considers at length involves a TV evangelist who, in the course of one of his sermons, says 'Jesus loves you.' Egan says:

What seems to be happening here is what we might think of as a sort of *shotgun assertion*, in which different asserted contents are going out to different audience members, rather than a single content going out to all of them. Each audience member gets their own assertion-pellet, loaded with its own proprietary content. (Egan (forthcoming): 17)

The metaphor of a shotgun assertion nicely captures the behind content relativism, and Egan's shows that a strong case can be made for a content relativistic treatment of the relevant data.

Illustration 4: Predicates Used in Instructions, Laws, and Imperatives

This last set of illustration involves some cases that have not so far been considered in this literature. The easiest way to see what I have in mind is to consider predicates used in certain kinds of instructions, orders, and laws. I focus here on instructions, but it should be obvious how to extend what I am saying to laws (and to various other aspects of the legal system). Imagine the following instructions given to airport employees:

- (1) Move suspicious packages away from crowds.
- (2) If a **dangerous** object is found in carry-on luggage, immediately contact a security officer.

I'll focus on the predicates 'suspicious' and 'dangerous'. Four points help bring out why content relativism can provide a natural account of their correct interpretation in (1) and (2):

• The instructor (i.e. the speaker of (1) and (2)) might realize that she is not in a position to anticipate all that should count as suspicious and dangerous in all relevant contexts of interpretation. To take this into account, she could intend for her instructions to be interpreted in ways that are best modelled by some version of content relativism. On this construal, the key predicates, 'suspicious' in (1) and 'dangerous' in (2), should be interpreted as varying their extensions from one interpreter to another.

This view of instructions should be familiar from the philosophy of law. A controversial but fairly mainstream view in legal theory is that legislators should be understood as intending for laws and other legal texts to have varying interpretation over time. According to such views, the correct interpretation of a law can change as social conditions change and as precedents develop over time.²⁰ The proposal is that this view of legal interpretation can provide a model of how certain instructions should be interpreted.

• So construed, (1) and (2) illustrate how there could be a *naturally restricted range of admissible interpreters*: those in a position to carry out the instructions. Again, a comparison to laws is illuminating. In the case of laws, there's a range of privileged interpreters and even a hierarchy of admissible interpreters (a higher court can override the interpretation offered by a lower court). Of course, we are all in some sense free to interpret laws (it's essential to a functional legal system that we do), but our interpretations have to defer to those of the privileged interpreters can be similarly restricted to those who are the intended recipients of the instructions (i.e., those being instructed). What you and I think of as dangerous is, in an important sense, overridden by the interpretation given by the intended audience, i.e. by those being instructed by (2).

- There is no guarantee that the content of the instructions, relative to an arbitrary admissible context of interpretation, is within the instructor's cognitive reach. She might have given instructions, relative to a certain context of interpretation, the exact content of which she could not anticipate at the time of speaking (in much the way that the correct interpretation of a law might not be restricted to contents the legislators had 'in mind' when the law was written.)
- Finally, it's not hard to see that norms can govern the act of giving instructions even when those instructions are understood as interpretation sensitive. The instructor can, for example, be held *responsible* for the contents she has expressed relative to various context of interpretation, even if she had no direct cognitive access to those contents. This point will be pursued in Section Four below.

I don't think content relativism provides the *only* model for how to understand these cases. There are other options worth exploring. The two most salient alternatives are:

- *Big character*: the instructor has to have in mind the extension of the relevant predicates for all contexts of interpretation. On this view, there are some very general conditions fixed in the context of utterance and these suffice to fix the relevant extensions relative to each context of interpretation (and so the content is the same across contexts of interpretation).
- *Utterance pluralism*: Egan (forthcoming) shows that one alternative to content relativism is to think of the original act as generating a plurality of new utterances, one for each interpreter. On this view, we shouldn't say that one utterance of an instruction has different contents relative to different contexts of interpretation. Rather, there's an event, a tokening of the sentence, which generates a range of other utterances. Each of these can have different contents (for more on this view, see Section Three below).

There might be utterances that are best explained by one of these alternative models. What I claim is only that some utterances of, e.g., (1) and (2) are best understood in an interpretation sensitive way. As I see it, the interpretation sensitive model is plausible in this case primarily because it provides a good model of the speaker's intentions. In the kinds of cases I have in mind, speakers intend for the contents of their utterances to be fixed in various contexts of interpretation other than the context of utterance (i.e., they don't intend for the context of utterance to be the only context of interpretation). If there are cases where it is plausible to describe the speakers as having such intentions, if having such intentions could be shown to have a clear communicative function, we have evidence for interpretation sensitivity theories over alternative models. Remember, I am not here taking a stand on the question of whether this level of content is properly labelled 'semantic content' (or on the question of whether it matters whether it is so labelled). What matters is only that there is some level of content that is interpretation sensitive. The claim is that in cases such as (1) and (2), the speaker intends that her utterances be interpretation sensitive, and such communicative intentions are, by any standard, important to a theory of communication.

Why would a speaker of (1) and (2) intend for her instructions to be interpretation sensitive? The reason is much the same as why a legislator might intend for their laws to be interpreted in different ways over time: she knows that her

instructions are to be implemented in contexts that have significant features that she cannot fully anticipate when she's giving the instructions. From a functional point of view, it makes sense for her to intend (and to get her audience to recognize that she intends) that key terms should have their extensions fixed in the context of interpretation, rather than the context of utterance.²¹

3. Comparisons

The reader should be excused for thinking that the characterization of content relativism so far leaves it unclear, to say the least, how to distinguish it from some closely related, but significantly different, views. On some construals, it is simply a version of the view Cappelen and Lepore (2004) call speech act pluralism. It is also a very close relative of utterance pluralism, briefly described above. Finally, on some construals, it is difficult to distinguish it from views that appeal to contextual vagueness. I discuss these in turn.

Comparison-1: Speech Act Pluralism

Cappelen and Lepore (1997, 2003, 2004) argue in favour of what they call 'Two-Dimensional Contextualism' (the two dimensions being context of utterance and context of interpretation) and against what they call 'Original Utterance Centrism' (the view that only the context of utterance, and not the context of interpretation, fixes what is said by the utterance of a sentence). In so doing, they argue for interpretation sensitivity broadly construed. However, another important component of the view defended in Cappelen and Lepore (2004), what they call 'speech act pluralism', stands in a more complex relationship to content relativism, and understanding this relationship is important (and difficult).

According to speech act pluralism, an assertive utterance of a declarative sentence can assert a plurality of propositions. Note, first, that this view need not imply content relativism. If the set of propositions asserted is the same relative to all contexts of interpretation, it does not. Of course, one could combine speech act pluralism with content relativism and the result would be a kind of view where the set of asserted propositions varies across contexts of interpretation.

So construed, the relationship between content relativism and speech act pluralism seems simple. To see why things aren't really that simple, consider the following challenge to the content relativist. Let u be an utterance by Mia of an interpretation sensitive sentence and let A and B be interpreters of u. According to the

²¹ In this section I have talked about certain *uses* of predicates being interpretation sensitive. Interpretation sensitivity is, however, a property of expressions, not of uses of expressions. Interpretation sensitivity applied to speech acts should be understood as an abbreviated way of talking about the interpretation sensitivity of the expressions used to perform the speech act. Predicates are interpretation sensitive, i.e. they have interpretation sensitive characters. Examples involving instructions are useful because in those cases it is particularly easy to see their interpretation sensitivity. Remember, as I construe interpretation sensitive interpretation sensitive sentence. In some cases, the range is restricted to the context of utterance (i.e., the context of utterance is the only admissible context of interpretations is restricted to the context of utterance. I don't take a stand on that issue here (in fact I doubt that it's true), but I focus on instructions because the phenomenon of interpretation sensitivity is easier to see in those cases.

content relativist, A can say truly, 'Mia said that P' and B can say truly, 'Mia said that Q'. Now, suppose that A and B hear each other's speech reports and that each knows that the other spoke truly. Shouldn't it then be true for A to say, based on the truth of B's speech report, that Mia said that Q, and for B to say, based on the truth of A's speech report, that Mia said that P? If so, they can both say truly that Mia said that P and Q. If so, there seems to be nothing left of the idea that one content is expressed relative to A and another relative to B. It looks, in effect, like content relativism has collapsed into speech act pluralism.²²

Maybe the most natural move for the content relativist in response to this challenge is to claim that 'said' (and related terms, like 'assert', 'claim', etc.) are context sensitive and have a hidden argument place that takes interpreters as values. If so, A's utterance of, 'Mia said that P' expresses the proposition *that Mia said_{for A} that* P, and B's utterance of that sentence expresses the proposition *that Mia said_{for B} that* Q. The content relativist's claim is that while these are both true, it is not true *that Mia said_{for B} that P* or *that Mia said_{for A} that Q*.

Note that this would not immediately block A from saying something true by uttering, 'Mia said that Q'. A could do so if she could get the hidden argument place in 'said' to take B as its value (sort of in the way we can get the location in 'It's raining' to shift away from the location of the speaker.) If A succeeds in this, the proposition she expresses by uttering 'Mia said that Q' is the proposition *that Mia* said_{forB} that Q and this is true (and compatible with the claim that Mia didn't say_{forA} that Q.)

This is one strategy content relativists can use to prevent collapse into speech act pluralism. There might be other such strategies, but I'm not clear on what they would be. Note that postulating an argument place for an interpreter in 'said' is not a trivial commitment. It's very much in need of independent justification. It goes beyond the scope of this paper to explore ways of providing such justifications. (See Cappelen and Lepore (1997) for some thoughts on these issues.)

Comparison-2: Utterance Pluralism

Whereas content relativists about a sentence S hold that an utterance of S can express different contents relative to different contexts of interpretation, an utterance pluralist holds that an 'utterance' (or, better, 'proto-utterance') of S in a context C can be used to generate different utterances relative to different contexts of interpretation. So, whereas the content relativist treats the speaker as having produced one utterance, in the context of utterance, that has different contents relative to different contexts of interpretation, the utterance pluralist thinks of that speaker as having produced a proto-utterance. Proto-utterances are, somehow, capable of generating a range of further utterances (one for each interpreter), each of which has just one content.

How can we adjudicate between content relativism and utterance pluralism? I'll restrict my discussion of the difference between the views to what I take to be one particularly important point:²³ the instances of interpretation sensitivity I find most plausible all allow for what I above called *mixed interpretations*, and content relativists have an easier time dealing with these cases. Here is a natural way to describe what happens in mixed interpretations: one part of the sentences has its content fixed in the context of utterance and another part of the sentences has its

²² There's still room for a weak form of content relativism, see Cappelen (2008).

²³ For more further discussion of the relationship between the views, see Egan (forthcoming).

content fixed in the context of interpretation. If that's the correct description, we need to appeal both to contexts of utterance and contexts of interpretation. The utterance pluralist only has access to what she calls 'contexts of utterance', and so can't explain account for this phenomenon. As an illustration, suppose epistemic modals are interpretation sensitive. When interpreting an utterance made in C_1 , of 'My uncle might be a crook' in a context of interpretation C_2 , the semantic value of 'my uncle' should be fixed in C_1 , while the semantic value of the epistemic modal is fixed in C_2 . A view that treats the context of interpretation as a context of utterance is deprived of this flexibility. These kinds of considerations seem both to distinguish content relativism from utterance pluralism and give the former an edge over the latter.

There are ways an utterance pluralist could try to account for mixed interpretations. Here is one option. Call the context of the proto-utterance (where this, roughly speaking, is the context in which the sentence's token is produced²⁴) the 'P-context'. According to the utterance pluralist, a proto-utterance generates a plurality of utterances. Each of these utterances takes place in separate contexts of utterance, $C_1...C_n$. The utterance pluralist could say that each of $C_1...C_n$ takes some of its values from the P-context. For example, the value of the speaker index in each of $C_1...C_n$ could automatically be fixed as the speaker in the P-context. The hope would be that the material needed to account for mixed interpretations is available in each of $C_1...C_n$, by virtue of the contextual parameters carried over, so to speak, from the P-context to each of the utterance contexts.

Whether such a strategy can succeed depends in large part on the details of particular proposals, and I won't try to settle the issue here. I will just note my suspicion that if an utterance pluralist can succeed in so construing her contexts of utterance, it will be hard, if not impossible, to tell content relativism and utterance pluralism apart. They will, I predict, end up being notational variations of the same view.²⁵

Comparison-3: Contextual IndeterminacyVagueness

Content relativism should be distinguished from a view according to which it is vague or indeterminate which context a speaker is in.²⁶ The case of epistemic modals can be used to illustrate the difference. Suppose the content of a sentence containing an epistemic modal includes a reference to the information state of a salient group or individual. Suppose further that this salient individual or group is fixed in the context of utterance. Finally, suppose that it could be indeterminate who should be considered the salient individual, or, more likely, who should be considered members of the salient group. One way to describe such a situation is to say that it is indeterminate (or vague) which context the sentence is uttered in. There are many candidate contexts of utterance (each one determining a particular individual or group) and it is indeterminate which context the utterance took place in.

²⁴ I note that this notion is not unproblematic: in some cases what counts as the production of the relevant token will be tricky (e.g., answering-machine-type cases). Cleaning up these issues will no doubt be a challenge for the utterance pluralist, but I won't pursue that concern here.

²⁵ In so far as they are different views, there seems to be some kind of metaphysical cost of having a lot of utterances generated by what, from a common sense point of view, is a single utterance. If utterance pluralism implies that I can generate a new utterance by a speaker just by remembering (or thinking about) one of her speech acts, that is, to put it mildly, a radical departure from our ordinary conception of an utterance.

²⁶ As mentioned above, von Fintel and Gillies describe their view as an appeal to contextual vagueness.

This phenomenon seems different from content relativism. There's no relativization of contents to contexts of interpretation and no need to distinguish between a context of utterance and a plurality of contexts of interpretation. However, if some assumptions are added to contextual vagueness, the view implies a version of content relativism. Suppose the speaker is aware of the contextual indeterminacy (she knows that all of $C_1...C_n$ are candidate contexts). Suppose we have a range of interpreters, I1...In and that each of them interprets the utterance relative to one of $C_1...C_n$. Suppose that it is also part of this view that all these interpreters are correct and that the speaker intends for all of them to be correct. We now have a view according to which two admissible interpreters, of the same utterance, get different interpretations, and both interpreters are right (and the speaker intends for both of the interpretations to be correct). This version of contextual vagueness is at least very close to what I have called content relativism above (the difference between the views might come down to nothing more than how the term 'utterance' should be used)

4. Content Relativism and Assertion

The central objection raised by Egan, Hawthorne, and Weatherson (2005) against content relativism is that they don't see how to reconcile it with their favoured theory of assertion. For reasons I'll make clear below, the content relativist can easily respond to these worries, but, more importantly, their argumentative strategy assumes that if interpretation sensitivity is hard to reconcile with standard conceptions of assertion, we should reject interpretation sensitivity. An equally appropriate response, I'll argue, is to revise standard conceptions of assertion.²⁷

There is little agreement about what an assertoric commitment amounts to. Some think assertion is a speech act governed by constitutive norms (Williamson (2000): Chapter 11). Those who hold this view don't agree on what the norm is.²⁸ Others think there is no stable norm for assertion (Levin,(forthcoming))—it varies from context to context. Others think assertion should be spelled out in non-normative terms.²⁹ For the sake of argument, I'll assume in what follows that belief is a norm for assertion, but not much in my argument depends on this choice – I use it simply to illustrate the kind of worry I have in mind.

Here, in abstract, is the problem. Suppose A utters S and, unbeknownst to A, this utterance means that P relative to some context of interpretation, C. Can we demand that A satisfy the norm of assertion with respect to P? Does that even make

²⁷ In this respect, I agree with an aspect of what John MacFarlane says in 'Making Sense of Relative Truth'. As MacFarlane sees it, the main challenge for any version of relativism about truth is to make sense of assertive commitment to an 'assessment sensitive proposition'. He concludes that paper by saying: 'We have given an account of assertoric commitment that settles just what one is committing oneself to in asserting an assessment-sensitive proposition. By doing this, I suggest, we have made relativism about truth intelligible.' (321). The challenge of giving an account of assertion in a content relativistic language is very different from doing so within a truth relativistic framework. Nonetheless, for reasons I make clear below, I think some of the strategies used by MacFarlane can be adopted by a content relativist.

²⁸ According to Timothy Williamson, You should: Assert P only if you know that P. According to others the norm is what it would be reasonable to believe (Lackey (2008)), or truth (Weiner (2005)).

²⁹ Others again, think the category of assertion is theoretically superfluous (Cappelen (forthcoming)). I'm in a bit of an awkward dialectical position here, since my own view of assertion is that it doesn't exist, that it's a philosophical illusion that we ever perform acts that should be classified as assertions. The argument in this section is directed at those who don't share this view. The goal is to show how they can make sense of assertion within a content relativistic framework.

sense? It is hard to see how she can commit (and audiences can take her as having made the commitment) to belief in a proposition that, in some significant sense, is outside her cognitive reach.

As an illustration, consider epistemic modals (not because I think this issue is particularly problematic for content relativism about epistemic modals, but because it provides a useful illustration). Amy utters, 'Max might be a crook'. Suppose that Lola is an admissible interpreter. Content relativism about epistemic modals predicts the truth of the following report:

Lola:	Amy said that, for all I know, Max is a crook.
Lola:	Amy said something about my epistemic state.

Assuming that belief is a norm for assertion:

Lola: Amy believes that, for all I know, Max is a crook.

These are strange consequences. After all, nothing in what I've said so far guarantees that Amy even knows who Lola is. How, then, can she have any beliefs about Lola's epistemic states? In what follows, I consider three replies to this worry, ranging from a very conservative reply to a very radical reply, and give some arguments in favour of the latter.

Reply One: Assertion Restricted to Proposition in Cu=Ci

First, consider a conservative reaction on behalf of content relativism according to which assertive commitment is only to the content expressed relative to the context of utterance considered as a context of interpretation. On this view, the speaker only asserts the content that results when Cu=Ci (I'll use 'Cu=Ci' as shorthand for 'the context of utterance is the context of interpretation'). On the view now under consideration, if S is interpretation sensitive, an utterance of S will vary in content across contexts of interpretation, but the speaker has only asserted *one* of these, the one we get when Cu=Ci.³⁰

This doesn't seem like an attractive reply on behalf of the content relativist. What is the point of describing P as the content of u relative to a content of interpretation C1 (where $Cu \neq C1$), if the speaker didn't assert P or say that P relative to C1? I don't want content relativists to be stuck with a view according to which interpreters in C1 can register that an utterance of S has P as its correct interpretation relative to C1, but cannot treat P as what was said or asserted by the speaker. If, in order to dis/agree with, criticize, or otherwise engage with what the speaker said, interpreters in C1have

³⁰ This might be the way to interpret Fintel and Gillies in the following passage:

We would like to point out that this "cloud of contexts" picture might also give us a handle on the norms of asserting might claims. The idea would be that when one asserts a sentence that is open to a variety of admissible interpretations, one needs to be justified in asserting it under the weakest admissible interpretation. In the case of might claims, that would typically mean that one has to be justified in asserting them as claims about one's own state of evidence. Of course, the hearer might interpret the sentence in a stronger fashion—and the speaker may well have anticipated and intended that—but the speaker can always retreat to the weakest interpretation, so that interpretation is the one she needs to be able to defend. (Fintel and Gillies (2008): 97.)

to appeal to what she said when Cu=Ci, then the claim that P is the correct interpretation in C1 seems to be nothing but a theoretical spinning wheel.

Reply Two: Safe Content Relativism

Proponents of content relativism who think that the range of assertion should be extended beyond cases where Cu=Ci, can allow *safe* or *risky* extensions. Risky versions of content relativism allow that speakers can say and assert contents that are, in some significant sense (still to be made precise), outside their cognitive reach. Safe versions don't allow this. So a reply slightly less conservative than the one considered above is to allow interpretation sensitive assertions as long as the range of admissible interpretations ensures that the asserted contents are within the speaker's cognitive reach.

This strategy, while open to the content relativist, has two potentially problematic features. First, it requires that we be able to spell out 'safe' and 'cognitively accessible' in a way that doesn't just rely on 'assertable', otherwise it would not be much help in resolving the above worry (we would in effect just be saying that the norm is that we can assert that P relative to Ci when P is assertable relative to Ci). Second—and this is the most important point—it is a motivated restriction on content relativism only if there's good reason for rejecting risky content relativism. I now move on to argue that there is no such reason.

Reply Three: Risky Content Relativism

Here are three reasons for not being worried about the idea that a speaker can assert contents that are cognitively inaccessible to her, i.e., for not being worried about risky assertions. The first point is just a reminder that most of us are committed to the assertion of such contents independently of any considerations having to do with interpretation sensitivity. The next two provide some suggestions for how content relativists could adjust our standard conceptions of assertion.

1. Those worried about risky content relativism should bear in mind that there is nothing surprising about asserted contents being fixed by factors that in some important sense lie outside the speaker's cognitive reach. Here are two illustrations of this point. First, on all standard conceptions of what contexts of utterances are, their content-fixing features can generate contents that are outside the speaker's cognitive reach. We can, for example, use 'here', 'now', 'you', and 'that' without knowing where we are, what time it is, whom we're speaking to, or what we are demonstrating. The more complex your conception of what a context is and the more complicated your meta-semantics, the more likely it will be that our utterances can express contents that in some significant sense are cognitively inaccessible us. Second, those of us who are externalists about contents are already committed to the idea that the contents we express by sentences containing names and natural kind terms can be cognitively inaccessible to us. According to such theories, we can use names and natural kind terms to talk and think about referents that, in some important sense, can lie beyond our cognitive reach. If you don't think externalism of this kind is a barrier to assertion, you have no good reason to be worried about risky content relativistic assertions. More generally, in so far as risky content relativism makes possible the assertion of cognitively inaccessible contents, this doesn't introduce any *new* puzzles about assertion.

2. Content relativism allows for a kind of flexibility in assertoric norms that can alleviate concerns about risky assertions. Suppose N is your favoured norm of assertion (i.e. belief, knowledge, truth or what have you). If you are a content relativist, you can let N govern assertion relative to Cu=Ci. Then let the norm for assertion vary across other contexts of interpretation. Here is an analogy. Some think it implausible that there is one norm governing all assertions. They think the norm of assertion is flexible – it varies across contexts of utterance.³¹ My proposal on behalf of the content relativist is analogous: even if you insist on there being a single norm that in some sense governs assertion in *any context of utterance taken as a context of interpretation*, you can hold that the norm relative to other contexts of interpretation varies. The epistemic constraints, for example, will typically be less severe with respect to an assertion of a content that's not entirely cognitively accessible.

So far, I've talked loosely of variable norms and their relative strength and weakness. What kind of norms might govern assertions of contents that have a certain degree of cognitive inaccessibility? I'll provide no conclusive answer to this question here, but it's useful to think about a suggestion inspired by the work of John MacFarlane³². The norms³³ MacFarlane proposes on behalf of the truth relativist can be adjusted to serve the content relativist's purposes. Let u be an assertive utterance of S by A and let P be the content of u relative to an arbitrary admissible context of interpretation Ci. Here is a rough MacFarlane-inspired description of some kinds of commitments A might make (note that these don't require that A believe what she asserts):

- (W) A is committed to withdrawing P if P is shown to be untrue.
- (J) A is committed to justifying P when it is appropriately challenged.
- (R) A is committed to accept responsibility if on the basis of this assertion someone else takes P to be true and it proves to be untrue.

These are illustrations of the kinds of commitments speakers of interpretation sensitive sentences can take on, and they don't require that the speaker believe the contents so expressed.³⁴ What I want to emphasize here—and this diverges radically from how MacFarlane (2005) uses the truth relativistic versions of these norms—is that endorsement of these as norms for contents expressed in cognitively inaccessible contexts of interpretation is compatible with some other norm(s) being constitutive of assertion when Cu=Ci. Suppose N is your favourite norm of assertion (you might even think it is essential to the speech act of assertion that N govern it.) You can let N be the norm for assertion for Cu=Ci, and let something like (W), (J), and (R) be the norms that govern contents expressed relative to other contexts of interpretation.

This still leaves open the question of whether it would be rational to make (W), (J), and (R) commitments to contents that are outside our cognitive reach. The answer to that question depends significantly on how serious that kind of commitment

³¹ See for example Levin (forthcoming).

³² See MacFarlane (2005).

 ³³ MacFarlane, I should mention, doesn't thinks of these as norms, but I will treat them here as such.
 ³⁴ The proposals raise a number of tricky questions: are these commitments to what A would do 'in Ci', and, if so, what is it for A to be in Ci? I won't try to address these questions here.

is. In that connection it is important to note that even on standard conceptions of assertion (e.g., Williamson's), the commitments associated with assertion are often cheap, in the sense that violations are ubiquitous and often cost-free. In the next paragraphs I suggest that proponents of risky content relativism should rely heavily on that point.

3. Proponents of risky content relativism should emphasise that those who endorse various norms for assertion seem to be comfortable with ubiquitous violation of these norms. The requirement, according to Williamson, is that we see ourselves as in some sense governed by the norm (i.e., we recognize the norm as governing our behaviour). Williamson, for example, holds that we should only assert what we know, even though he recognizes that speakers very often don't stand in the knowledge relation to the propositions they assert. He doesn't think this is a serious challenge to his view. We often engage in rule-governed behaviour, he says, even though the rules governing us are regularly broken without any kind of sanction. In the case of assertion, ubiquitous rule-braking just shows that assertions often come cheap. Williamson says: 'When assertions come cheap, it is not because the knowledge rule is no longer in force, but because violations of the rule have ceased to matter so much'. (Williamson (2000): 259)

Think of risky content relativism as introducing another source of cheap assertions: we're in some sense governed by various norms (e.g., those described by (W), (J) and (R)), but their violations don't matter much. When we assertively utter interpretation sensitive sentences we make commitments to a range of contents, in a range of contexts of interpretation. These commitments are often cheap, to use Williamson's terminology. But that should come as no surprise to those who think of assertion as governed by norms.

I conclude that risky content relativism is compatible with the idea that assertion is a norm-governed activity (it's even compatible with the view that such norms are constitutive of assertion). There is, however, another strand to the literature on assertion that I have not so far addressed. According to Robert Stalnaker, the 'essential effect of assertion' (Stalnaker (1978)) is to add the content expressed to the context set. Egan, Hawthorne, and Weatherson (2005) argue that content relativism is incompatible with this view of assertion, and I next turn to that criticism. Replying to this objection is fairly straightforward given what I've said about norms of assertion above.

Content Relativism and Stalnaker on the 'Essential Effect of Assertion'

Here is one of Egan, Hawthorne, and Weatherson's central objections to content relativism:

[...] it causes difficulties for an attractive part of the Stalnakerian story about assertion, that the central role of an assertion is to add the proposition asserted to the stock of conversational presuppositions (Stalnaker 1978). On the content relativist view, it can't be that the essential effect of assertion is to add *the proposition asserted* to the stock of common presuppositions, because there's no such thing as *the* proposition asserted.

Given what I've said about norms of assertion above, it should be clear why this

presents no serious challenge to the content relativist. There are several options, two of which are these:

Option 1: The content relativist can treat Stalnaker's claim as applying to the content asserted when Cu=Ci. On this construal, Stalnaker says that it is an essential feature of assertion that if P is expressed in Cu=Ci, then P is added to the context set.

Option 2: She can let Stalnaker's claim apply to all contexts of interpretation, i.e., for any context of interpretation Ci, the content expressed relative to Ci is added to Ci's context set.

There are intermediate positions available. She can, for example, limit the Stalnakerian story to cognitively accessible contexts of interpretation (that would be the analogue of the intermediate, safe option outlined above).

Considerations analogous to those discussed in connection with the norms of assertion will be relevant when evaluating these options. With respect to Option 1, it's reasonable to ask: if, for a context of interpretation, Ci, the proposition expressed by u, say, P, isn't added to the context set of Ci, what's the point of saying that P is asserted relative to Ci? As in the case of the more conservative reply above, the threat is that content relativism might turn into a kind of theoretical spinning wheel, with no real function. With respect to Option 2, we can ask: how can the speaker intend to add P to the context set of a context of interpretation, Ci, if P and Ci are cognitively inaccessible her?

I won't try here to adjudicate between these options, but it should be clear how the arguments given for risky versions of content relativism above can be used in favour of Option 2. There are, however, additional complications in this case. What one says here will depend on how one understands Stalnaker's view more generally, how one thinks of context sets, and what theoretical role these play. The correct story might also vary across different kinds of expressions. What should be clear, however, is that, if risky content relativistic assertions are acceptable for reasons given above, then no knock-down objections to content relativism are likely to be forthcoming from these kinds of considerations. There might be interesting things to say and I suspect that they will illuminate and enrich the Stalnakerian story about assertion, but we have been given no reason to think that any of this will undermine content relativism.

Conclusion

From the point of view of mainstream philosophy of language, the idea that interpreters can be co-creators of contents will seem both bizarre and, hopefully, titillatingly radical.³⁵ It's a view that contradicts many mainstream paradigms for theorizing about communication. We should, therefore, endorse content relativism only with a great deal of caution. In this paper I've covered some, but certainly not all, of the terrain that needs to be navigated in order to develop an acceptable version of the view. A first important step towards legitimizing the framework is to rule out principled objections to the view – to show that there are no master arguments to the

³⁵ One theory of communication that incorporate some components of content relativism is relevance theory; see e.g. Sperber and Wilson (1986) and Robyn Carston (2002).

effect that *any* attempt to implement the view is doomed to failure. The best candidates for such objections are, I think, those having to do with assertion. For reasons spelled out above, those arguments are utterly unconvincing; when properly understood, content relativism will improve our understanding of the nature of assertion, not undermine it.

If we put aside the thought that there's some principled objection to content relativism, what remains is still hard work: to show, in particular cases, that our theories are improved by thinking of classes of sentences as exhibiting sensitivity not only to a context of utterance, but also to contexts of interpretation. The closest we got to a very general argument for this view is the potential that instructions, imperatives, and laws (and other legal texts) have to exhibit interpretation sensitivity. If that turns out to be our best explanation of what goes on in those cases, we would seem to have a very general argument for interpretation sensitivity, since the examples involved are not limited to any sub-set of expressions—the arguments apply to expressions as used in the performance of certain kinds of speech acts, and there's no restriction on the kinds of expression that can be so used.

REFERENCES

Bach, K., (1994) 'Conversational Impliciture', Mind and Language 9: 124-162

- Cappelen, H. (forthcoming). 'Against Assertion' in Brown, J. and Cappelen, H. (eds.) New *Essays on Assertion.* Oxford: Oxford University Press.
- Cappelen, H. (2008). 'Content Relativism', in Garcia-Carpintero, M. and Kölbel, M. (eds.) *Relative Truth*. Oxford: Oxford University Press.
- Cappelen, H. (2007). 'Semantics and Pragmatics: Some Central Issues' in Preyer, G. and Peter, G. (eds.), *Introduction to Context Sensitivity and Semantic Minimalism: New Essays on Semantics and Pragmatics*. Oxford: Oxford University Press.
- Cappelen, H. and Hawthorne, J. (forthcoming). *Relativism and Monadic Truth*. Oxford: Oxford University Press.
- Cappelen, H. and Lepore, E. (2004). *Insensitive Semantics: A Defense of Minimalism and Speech Act Pluralism.* Malden, MA: Blackwell.
- Cappelen, H. and Lepore, E. (2003). 'Context Shifting Arguments', in Hawthorne, J., and Zimmerman, D. (eds), *Philosophical Perspectives 17*. Language and *Philosophical Linguistics*. Malden, MA: Blackwell: 25-50.
- Cappelen, H. and Lepore. E. (1997). 'On an Alleged Connection Between Indirect Speech and The Theory of Meaning' *Mind and Language* 12: 278-296.
- Carston, R. (2002). *Thoughts and Utterances: The Pragmatics of Explicit Communication*. Malden, Ma.: Blackwell.
- Endicott, T. (2001). Vagueness and Law, Oxford University Press.

- Egan, A. (forthcoming). 'Billboards, Bombs, and Shotgun Weddings', *Synthese*. Available at: <u>http://sitemaker.umich.edu/egana/files/bilboards.2007.06.22.pdf</u> Page references are to the online version.
- Egan, A., Hawthorne, J., and Weatherson, B. (2005) 'Epistemic Modals and Relative Truth', in Peter, G., and Preyer, G. (eds.), *Contextualism in Philosophy*. Oxford: Clarendon: 131-169.

von Fintel, K. and Gillies, A. (2008). 'CIA leaks', Philosophical Review 117(1): 77-98.

- Hart, H.L.A. The Concept of Law, Oxford University Press
- Kaplan, D. (1989). 'Demonstratives', in Almog, J., Perry, J., and Wettstein, H. (eds.), *Themes from Kaplan*. Oxford: Oxford University Press: 481-563.
- Lackey, J. (2008). *Learning from Words: Testimony as a Source of Knowledge*. Oxford: Oxford University Press.
- Lasersohn, P. (2005). 'Context Dependence, Disagreement and Predicates of Personal Taste', *Linguistics and Philosophy* 28(6): 643-686.
- Levin, J. (forthcoming). 'Assertion, Practical Reason, and Pragmatic Theories of Knowledge', *Philosophy and Phenomenological Research.*
- MacFarlane, J. (2005), 'Making Sense of Relative Truth', *Proceedings of the Aristotelian* Society 105:1: 305-323.
- Salmon, N. (1991). 'The Pragmatic Fallacy', Philosophical Studies 63: 83-97.
- Soames, S. (2002). *Beyond Rigidity: The Unfinished Agenda of Naming and Necessity*. Oxford: Oxford University Press.
- Sperber, D., and Wilson, D. (1986). *Relevance: Communication and Cognition*. Oxford: Blackwell.
- Stephenson, T. (2007). 'Judge Dependence, Epistemic Modals, and Predicates of Personal Taste', *Linguistics and Philosophy* 30(4): 487-525.
- Stalnaker, R. (1978). 'Assertion', in Stalnaker, R. (ed.). Context and Content: Essays on Intentionality in Speech and Thought. Oxford: Oxford University Press, 78-95.
- Weatherson, B., (forthcoming). 'Conditionals and Indexical Relativism', available at <u>http://brian.weatherson.org/CaIR.pdf</u>
- Weiner, M. (2005). 'Must We Know What We Say?', Philosophical Review 114: 227-251.

Williamson, T. (2000). Knowledge and Its Limits. Oxford: Oxford University Press.