

## Agreement and Communication

Max Kölbel, LOGOS Research Group, ICREA at Universitat de Barcelona

Forthcoming in *Erkenntnis*. Version of 27 June 2012.

Abstract: I distinguish two notions of agreement (disagreement) in belief: (a) believing the same (contradictory) content(s) vs (b) having beliefs that necessarily coincide/diverge in normative status. The second notion of agreement (disagreement), (b), is clearly significant for the communication of beliefs amongst thinkers. Thus there would seem to be some *prima facie* advantage to choosing the conception of content operative in (a) in such a way that the normative status of beliefs supervenes on their content, and this seems to be the prevailing assumption of many semanticists. I shall argue that *de se* beliefs and assertions provide a motivation to depart from this assumption, and so do beliefs and assertions concerning what is epistemically possible. I conclude by offering two models of assertoric communication that are compatible with the abandonment of the assumption, and suggesting schematically that each model applies to different cases.

### 1. Two Notions of Agreement and Disagreement

Recently John MacFarlane and others have sought to give an “account of disagreement” (MacFarlane 2007, p. 24) “to capture the core of the idea of disagreement” (ibid, p. 25), or to “get clear about what it is for two people to disagree” (ibid, p. 18).<sup>1</sup> In these discussions, theorists have been trying to capture a notion of disagreement that conforms to intuitions or everyday judgements as to what it takes for two people “really” or “genuinely” to disagree. In this paper, I am *not* going to pursue this question: I am not going to try to say what it is genuinely to disagree. I am unsure as to whether there is any clear, univocal pretheoretical notion of disagreement that it makes sense to try to analyse or account for in this way. Nevertheless, what I want to say is relevant to these discussions. I will be saying something about two notions that one might with some justice call “notions of disagreement”, and about the significance of their relation to one another in semantics. However, I am not pursuing the question whether either of these notions of disagreement has a good claim on being *the* notion of disagreement or *genuine* disagreement. All I shall be interested in is the two notions I’ll be defining, without any claim as to the extent to, or manner in, which they correspond to *the* notion of genuine disagreement that figures in the recent debates, or *the* alleged notion of disagreement of our pretheoretical repertoire of concepts.

More precisely, I shall be talking about two notions of agreement and two corresponding notions of disagreement, both of them stipulatively defined. I shall begin with two senses of agreement and then specify the corresponding notions of disagreement.

Saying in general what agreement is is a bit like saying what equality or similarity is: agreement concerns a *respect* of agreement, just like equality and similarity concern respects of similarity or equality. For example, in grammar a noun and an adjective might agree in respect of gender, but not agree in respect of number. Two figures might be equal in shape but not in colour. Presumably, disagreement comes in

---

<sup>1</sup> Further examples include Stojanovic 2007, Cappelen & Hawthorne 2009 and Sundell 2011.

as many respects as agreement. For the purposes of this paper, we are only interested in agreement (and disagreement) amongst entities in a certain highly restricted domain: agreement or disagreement amongst thinkers in so far as they have a certain belief, or in so far as they are asserting something. For simplicity, we can speak of two beliefs or assertions agreeing or being in agreement—this is not how we normally speak, but that’s not a problem, as I am stipulating (see above). The point is to distinguish two notions of agreement amongst beliefs or assertions—or if you wish: agreement amongst thinkers in so far as they have these beliefs or make these assertions.

Agreement in my first stipulated sense is *coincidence of content*. Contents (propositions) are theoretical abstract entities, postulated in semantics, and perhaps in theories of mental representation. A variety of different kinds of entities have been given the role of contents in theories of representation, and contents have served a variety of different purposes in such theories. At a minimum, the assignment of contents serves to explain some of the logical relations amongst the representations to which they are assigned. To fix ideas further, let us say that we are operating in a framework where speakers *assert* contents by using sentences, and that these assertions count as *sincere* just if the asserter believes the very same content he or she asserted. Furthermore, let’s assume that a semantic theory assigns to each sentence the semantic content it expresses in a context of utterance, and that using a declarative sentence in a context counts as asserting precisely the proposition that is the semantic content of the sentence at that context. For easy reference, let me list these framework assumptions:

- A1 There are abstract entities—called “contents” or “propositions”—which are used to characterise both beliefs and assertions, by saying that they are the contents “of” these beliefs or assertions. (In other words: beliefs are propositional attitudes and assertions are propositional acts.)
- A2 The assertion of a content is sincere if and only if the asserter believes the content asserted.<sup>2</sup>
- A3 The semantics of a natural language assigns “semantic contents” to sentences in context. When a sentence is used to make an assertion, it is the semantic content of the sentence in the context in which it was used which is the content of the assertion. (Short: semantic content = assertoric content.)<sup>3</sup>

---

<sup>2</sup> Two comments: 1. This assumption will not be shared, for example, by neo-Russellians about names, such as Perry (1977, 1979), Salmon (1986), Soames (2002) or Nelson (2002). Thus, the framework I am assuming here excludes these accounts. However, the possibility of dropping assumption A2 will be considered explicitly below. 2. It might be thought that believing the content asserted is not sufficient for sincerity because one’s assertion might be causally unconnected to the belief and therefore not be sincere. However, it is doubtful whether credible such cases will arise independently of an opaqueness of content thesis. Thanks to Bryan Pickel and an anonymous referee for raising the issue.

<sup>3</sup> The question of how the semantic contents of sentences are related to the contents of assertions that are performed by uttering these sentences is in fact one that has generated some discussion (see Lewis 1980, Dummett 1991, Stanley 1997, Egan 2007, Moss forthcoming, Rabern forthcoming a and b). The very simple assumption made here (i.e. the relation is that of identity) is just the simplest extension of the framework in Kaplan 1977/1989, see *ibid.* p. 522. I will be revisiting this point below, i.e. I will consider dropping assumption A3.

The first kind of agreement, then, is agreement in content. Let's call this "agreementC", where the suffixed "C" is supposed to remind us of "coincidence of content". Here is a definition:

(C) Two beliefs or assertions agreeC if and only if they have the same content.

The second notion of agreement concerns the norms to which beliefs and assertions are subject. Explaining this notion will require some further setting up. Let us start with beliefs. Beliefs are subject to norms of correct believing, in the sense of "correct" according to which a belief is correct if things are as it represents them as being. (I am avoiding the word "true" merely in order not to interfere with any philosophical views about truth.) In my framework of thinking, what it is correct to believe depends on the concepts involved in what is believed. Thus, the contents of belief (propositions, contents) will be constituted by concepts, and concepts come with norms of correct application. Concepts such as that of redness have an extension: the range of objects to which the concept can be correctly applied. In order not to rule out any views, we could say more generally that it is constitutive of each concept under which conditions it is correct to apply it. I am deliberately not saying that each concept has an extension of objects to which it is correct to apply it, because I want to leave open the possibility that the extension of a concept not be the same for every thinker, or one thinker at any time.

What determines which concepts are involved in a given belief? In the framework I am using here<sup>4</sup>, the idea is that thinkers come to use concepts by being a member of a community of thinkers, often the same community within which that thinker has acquired their own conceptual competence. Such a community will employ practices of mutual calibration and correction, thus making sure that members' thinking efforts are co-ordinated, in that common methods are used for establishing whether a concept is correctly applied in a given case etc. The concepts used by the thinkers are those of which an idealised version of their thinking practices would make most sense.<sup>5</sup>

Being a competent thinker in a community does not guarantee that one has only correct beliefs. For even a generally competent thinker will occasionally make mistakes (even make *frequent* mistakes, if circumstances are unfavourable). Thus, there is a difference between competence and performance for believers. What I am interested in here is correctness in the sense of performance in line with competence. That is, I am assuming that the correctness of a performance consist in compliance with certain norms governing the concepts employed, norms which arise from certain practices of mutual calibration and correction in the community of which the thinker is a member. There are many further interesting questions about this framework, but I believe that its key aspects for the purposes of this article have been stated: there are norms regarding the correct application of concepts, and the correctness of beliefs depends on these conceptual norms.

Derivatively, an *assertion* can be correct in the sense that it would be correct for the asserter to believe the content asserted.

---

<sup>4</sup> This framework is described in some more detail in Kölbel forthcoming.

<sup>5</sup> This not only makes room for concepts that vary in their extension from one thinker, time, place to another (which is relevant for our purposes here) but also for concepts that have vague extensions in the sense that they vary with a range of permissible sharpenings, none of which is privileged by the practices of the community in question. For discussion of this, see Kölbel 2010.

Now, the second notion of agreement I am interested in is the relationship of necessary coincidence in normative status. Thus, two beliefs (assertions) agree in this sense just if correctness of the first necessitates the correctness of the second and vice versa. The necessity in question is meant to be conceptual necessity, for the coincidence in normative status is supposed to hold as a result of conceptual norms, i.e. the norms of correct belief. Let's call this second form of agreement "agreement<sub>N</sub>", where "N" is supposed to remind us of "normative coincidence":

N Two beliefs or assertions agree<sub>N</sub> just if: necessarily, if the one is correct then so is the other, and vice versa.

It is easy to supplement our two definitions of agreement with corresponding definitions of disagreement. Just as two beliefs or assertions can have the same content, they can have contradictory contents. Accordingly, we can say that two beliefs or assertions disagree<sub>C</sub> when they have contradictory contents. What exactly is it for contents to contradict one another?

I am here interested in a semantic notion of contradiction: roughly, contradictory contents cannot both be true at once. But this also introduces a new idea to the framework, namely the idea that contents can have truth-values. Now, I want my framework to be neutral with regards to the question whether contents have absolute truth-values, whether the truth-value of a proposition depends on anything other than how things are in the (objective) world. For one of the questions that will occupy us below is the question of whether, for example, the very same content of thought and speech can be true relative to one location and not true relative to another. In order, therefore, to remain neutral, I shall merely assume that contents have truth-values relative to a "point of evaluation", leaving open what such a point of evaluation is. Such a point of evaluation might, or might not, involve factors that go beyond the possible state of the world. Let us add this to our framework assumptions:

A4 Contents have truth-values relative to points of evaluation. Points of evaluation involve a possible world (or the possible state of the world) as one, but not necessarily the only, factor.<sup>6</sup>

We can then define contradictoriness as follows: a content  $p$  contradicts a content  $q$  just if  $p$ 's truth at any point of evaluation entails that  $q$  is not true at that point of evaluation, and vice versa. In turn, we can define disagreement<sub>C</sub> as follows:

DC Two beliefs or assertions disagree<sub>C</sub> just if they have contradictory contents.

Similarly, just as beliefs or assertions can necessarily coincide in normative status, they can necessarily diverge in normative status, which gives us the notion of disagreement<sub>N</sub>:

DN Two beliefs or assertions disagree<sub>N</sub> just if: necessarily, if the one is correct then the other is not.

---

<sup>6</sup> On a view like Jonathan Schaffer's "necessitarianism", propositions are individuated in such a way that they do not vary in truth-value with the possible state of the world. Thus, on a such a view, A4 is at best pointless because vacuously true. For why should we treat contents as having truth-values relative to a point of evaluation, if their truth-value does not vary from one point to another. See Schaffer forthcoming.

## **2. Desideratum: agreementC determines agreementN**

AgreementN and disagreementN are interesting notions from the point of view of understanding the pooling of cognitive resources amongst thinkers. Knowing which potential beliefs would agreeN with some belief that is likely correct may allow one actually to acquire such beliefs, beliefs that are as likely to be correct as the original belief. If one person makes a sincere assertion, thereby expressing a belief, then if anyone in response forms a new belief that agreesN with the one expressed in the assertion, then this newly formed belief will be correct as long as the original belief was correct.

DisagreementN is similarly significant: of two disagreeingN beliefs at most one can be correct. In certain circumstances, this will provide motivation for finding out which of the two is not correct, and how the error came about. If these disagreeingN beliefs are expressed in sincere assertions, then the disagreeingN parties might engage in an argument, i.e. exchange the reasons for their respective views, thus potentially unmasking the error that one of them must be committing.

In a large number of cases, we seem to be able to adjudicate agreementN and disagreementN relations amongst potential beliefs and assertions merely on the basis of our general conceptual and linguistic competence. For example, if you assert that Huxley wrote *Homage to Catalonia*, I know that if your assertion is correct, then if I were to believe that Huxley wrote *Homage to Catalonia* it would also be correct (in fact, it is not correct, but the agreementN relation still holds). Thus I know that your potential assertion and my potential belief agreeN with one another. If I didn't know this, then this would amount to a deficiency in my competence with the expressions and concepts involved. It therefore seems reasonable to presume that our ability to adjudicate agreementN and disagreementN in an a priori manner is due to the fact that agreementN relations amongst beliefs and assertions are *conceptually* linked to the contents of these beliefs and assertions. Moreover, it is independently plausible that the norms of correctness governing beliefs and assertions should spring from the contents of beliefs and assertions: after all, contents can bear truth-values, and plausibly the correctness of beliefs and assertions should be related to their truth. Many, myself included, would say that the semantic contents assigned to utterances by a semantic theory should predict the judgements of competent speakers as to literal correctness (compare for example Carnap 1955, Stanley 2007 for similar approaches). It is therefore worth exploring the possible links between the two forms of agreement and disagreement; or between agreementN and and content.

One very simple hypothesis would be that the correctness of a belief or assertion consists simply in the truth of its content, and the incorrectness of a belief or assertion consists in the lack of truth of its content. A weaker hypothesis, one that I want to explore, is the hypothesis that facts about agreementN or disagreementN amongst beliefs and assertions can be fully explained in terms of the contents of these beliefs and assertions. It would follow that facts about the correctness of beliefs or assertions, and therefore about agreementN and disagreementN amongst beliefs or assertions supervene on facts about the contents of these beliefs or assertions. Let's call a semantic theory that assigns belief- and assertoric contents in such a way that the hypothesis is true a "straight" semantic theory. On a straight semantics, then, any beliefs or assertions that agreeC would also have to agreeN. And any beliefs or

assertions that do not agree<sub>N</sub> could not be beliefs or assertions that agree<sub>C</sub>.<sup>7</sup> In other words, a straight semantics assigns only *portable* contents: if it is correct for anyone at any time to believe such a content, then it is correct for anyone at any time to believe it. To repeat, then, a *straight* semantics is one which guarantees H:

H Facts about agreement<sub>N</sub> and disagreement<sub>N</sub> amongst beliefs or assertions supervene on facts about the contents of these beliefs or assertions.

Agreement<sub>C</sub> and disagreement<sub>C</sub> depend on our choice of a theory of content. Which beliefs and assertions agree<sub>C</sub> or disagree<sub>C</sub> depends on how we assign contents to beliefs and assertions. Moreover, via our assumption A3, the assignment of contents to beliefs and assertions is linked to our choice of semantic theory for natural languages. For the semantic content assigned by the semantics to a sentence in a given context will be, according to A3, the content of any assertion made by uttering the sentence in such a context, and this will also be the content belief in which is required for such an assertion to be sincere, as stated in A2. Is straightness a desirable feature in a semantic theory? I believe it is at least *prima facie* desirable, for it is desirable that the assignment of contents to beliefs and assertions explain facts about the normative status of beliefs and assertions in the simplest possible way.

Let me illustrate with an example the kind of choices we might have between alternative theories of content. Suppose at 12 noon Greenwich Mean Time (GMT) in London in summer Alfred has a belief that he sincerely expresses by uttering “It’s 1pm.” Simultaneously Bert is in Barcelona and also has a belief that he sincerely expresses by uttering “It’s 1pm.” Let’s assume that the belief and assertion of each of them concerns their respective locations. Due to the difference in time-zone, then, only one of them can be right. In fact, Alfred happens to be right, while Bert is wrong (British summer time = GMT + 1, while Spanish summer time = GMT + 2). Do Alfred and Bert agree<sub>C</sub> in their respective beliefs?

This depends on how we construe the contents of their beliefs. Let’s consider two different theories of content, Theory 1 and Theory 2. According to Theory 1, they do agree<sub>C</sub>, Alfred and Bert believe and assert the same content (we might describe it as the content that it is 1pm at that time, the time at which they have these beliefs). Whether it is correct to believe this shared content, however, depends on the time-zone concerned. So while Alfred and Bert believe and assert the same content, their beliefs and assertions concern different time zones, so that, despite their agreement<sub>C</sub>, Alfred’s belief is correct while Bert’s is not.

According to Theory 2, Alfred and Bert do not agree<sub>C</sub>, they believe and assert different contents. Alfred believes what might be called “the content that it is 1pm at that time *in London* (or in the time zone in which London is located)”, while Bert believes what might be called “the content that it is 1pm at that time *in Barcelona* (or in the time zone in which Barcelona is located)”. It is this difference in the contents of their beliefs that accounts for why Alfred’s belief is correct and Bert’s is not.

---

<sup>7</sup> If we were working in a framework of “flat” contents, i.e. a theory according to which there is only one necessary content (the content that has the value true with respect to all points of evaluation), and only one contradictory content (the content that is false with respect to all points of evaluation), such as those of Stalnaker 1984 or Lewis 1986, then we might instead be exploring the hypothesis that agreement<sub>C</sub> and agreement<sub>N</sub> coincide (and this was in fact the focus of the talks from which this paper derives). However, in order to be general and include frameworks of structured propositions (such as the “metaphysical picture” of Kaplan 1977, or King 2003), I am exploring merely the hypothesis that agreement<sub>N</sub> supervenes on content, so that agreement<sub>C</sub> entails agreement<sub>N</sub>.

Both theories can be motivated, neither can be dismissed out of hand.<sup>8</sup> So this example shows how a theory of belief content may or may not assign different contents to beliefs that fail to agreeN (or even disagreeN, as in this case). Let's look again at Alfred's belief at 12 noon GMT that he expresses by saying "It's 1pm.", and compare it with another belief, namely Carl's belief one hour later, also concerning London, that he would express by saying "It was 1pm one hour ago.". Do these beliefs agreeC? Again, it depends on our theory of content.

Plausibly, according to Theory 1, Alfred and Carl believe different contents. Alfred believes what might be described as the content that it is 1pm, while Carl believes what might be described as the content that it was 1pm an hour earlier. Nevertheless, necessarily, if Alfred's belief is correct, then so is Carl's, and vice versa. They agreeN. This can be explained by the fact that their beliefs have the respective contents they have *together with* the fact that Carl's belief occurs an hour after Alfred's. So, again, Theory 1 requires facts that go beyond the mere content of these beliefs in order to explain whether they are correct.

According to Theory 2, perhaps Alfred and Carl believe the same content, so they agreeC. The content believed by both could be described as the content that it is 1pm at 12 noon GMT.

As we can see, whether two beliefs agreeC depends on the theory of content we adopt. On some choices of theory of content, agreementC entails agreementN, but on other choices of theory it doesn't. Theory 2 is compatible with the idea that agreementC entails agreementN, while Theory 1 is not. Theory 1 cannot explain facts about the correctness of beliefs merely by reference to the contents of these beliefs but rather requires further information about the respective times at which they are held.

Is straightness a desirable feature in a semantic theory? I believe straightness has at least two advantages. First, a theory of content that can by itself explain facts concerning the correctness of beliefs and assertions seems more powerful than one that requires information about beliefs and assertions that goes beyond their content. Secondly, such a theory of the contents of beliefs and assertions would also provide an elegantly simple account of linguistic communication and how information is shared amongst thinkers. If the normative status of a belief or assertion depends only on its content, then there is no danger that, if different people believe the same content, one of them is wrong while the other is right. There is no possibility that the same content might be correctly believable on one occasion but not correctly believable on another. Thus we can think of linguistic communication, when successful, as simply the transmission of content from one thinker to another. A source sincerely asserts a content, a content that she believes, and the recipient comes to believe that very same content as a result of witnessing the assertion.

### **3. The Limits of Straightness: De Se**

Straight semantic theories do not only have advantages. They also have limitations. In this section I want to explain some of these limitations, limitations that provide motivation to explore non-straight semantic theories.

---

<sup>8</sup> To provide two examples: The type of theory proposed by Recanati in his 2007 suggests that he would treat location dependence as in Theory 1, while the type of proposals made by Stanley (2000, 2007) suggests that he would tend to accept a theory like Theory 2.

Before discussing the limitations I am interested in, I want to depart on a brief historical digression. Frege was famously an adherent of straightness: he thought that the content of an utterance alone (“the thought expressed”) was to determine its truth-value. If we observe that the very same sentence, e.g. “I am Bismarck.” can be true on one occasion and not true on another, then this shows that the thought expressed on the two occasions is different. But since all the expressions used are the same, there must be a further, unpronounced, element in the sentence, which expresses a further element in the content, which then accounts for the difference between the thought Bismarck expresses, when he utters “I am Bismarck.” and the thought I express when I utter the same words. Since I can never utter “I am Bismarck.” in such a way that it expresses the thought Bismarck expresses when he utters it, and Bismarck can never produce the thought I express when I utter it, Frege concluded that the extra element in the thoughts in question was something expressible and graspable only by one person, respectively Bismarck and me.<sup>9</sup> Frege was led to accept “egocentric” contents (contents only thinkable and assertable by one person<sup>10</sup>) largely because of his commitment to straightness. As we shall see, some contemporary views are similarly driven to accept egocentric contents because of an adherence to straightness.

The first limitation concerns so-called *de se* beliefs and assertions. Let us consider two potential beliefs, the beliefs that I would express<sup>11</sup> if in a certain situation *S* I were to utter, respectively, the following two sentences:

D1 My pants are on fire.

D2 MK’s pants are on fire.

The situation *S* in question is a situation where I have been told over a public address system that MK’s pants are on fire, and I believe what I have been told. However, I do not remember that I am MK, or perhaps better: I do not realise that the person about whom it was said that their pants are on fire is me. It seems clear that I can sincerely utter D2 in *S*, because I believe that MK’s pants are on fire. This means, that in some sense I can be said to believe that my pants are on fire.<sup>12</sup> However, it seems equally clear that to utter D1 would be insincere. I do not have the belief that my pants—under that description—are on fire. Now, there are differences between a belief that would make my utterance D1 sincere and a belief that would make my utterance D2 sincere.<sup>13</sup> Let’s call these two potential beliefs my “D1-belief” and my

---

<sup>9</sup> See Frege 1918. For discussion, see Künne 1992.

<sup>10</sup> Given Frege’s theory of the concept of truth, namely that it, when applied to a thought, yields the same thought, perhaps this is false: I can assert and think the thought Bismarck expresses when he says “I am Bismarck.” by simply saying “What Bismarck asserted is true.” and thinking the thought this expresses. But this depends on Frege’s problematic theory of truth, which, incidentally, seems to be in conflict with Frege’s own criterion of thought-difference: if the same person can rationally accept thought *t*<sub>1</sub> and reject thought *t*<sub>2</sub>, then *t*<sub>1</sub> ≠ *t*<sub>2</sub> (see Frege 1892).

<sup>11</sup> The potential belief expressed by an utterance is the belief one would need to have in order for the assertion to be sincere.

<sup>12</sup> It is not always appreciated that a report by me “I believed my pants were on fire.”, or by someone else about me “He believes his pants are on fire.”, can be correct even if I am not aware at the relevant time that these pants are mine. Cf. Moss forthcoming and Stanley 2011, ch. 3.

<sup>13</sup> More accurately: some beliefs will make utterance D2 sincere, but not utterance D1—this is what the scenario is supposed to illustrate. Some might deny even this. They might insist that if utterance D2 were sincere, then so would utterance D1. But this would cause the bump to go up in a different place again. In our scenario, I might have forgotten that I am MK and have the D2-belief. Still, I might want to pretend to an audience that I am MK: “MK’s pants are on fire—that’s me! My pants are on fire!”. Utterance D1 here seems to be a lie. The theorist who claims that the D1-belief and the D2-belief have

“D2-belief” respectively. My D1-belief would typically cause alarm in me and move me to take self-saving measures. My D2-belief will at best cause a different kind of alarm, and it will not by itself motivate self-saving behaviour.<sup>14</sup>

A theory of content will either say that my D1-belief and my D2 belief have the same content or that they have distinct contents. The first possibility is ruled out by our assumptions A2 and A3. For if, as A3 says, the content of assertion is simply the semantic content of the expression uttered in the assertion, and the semantic content of the potential utterances of D1 and D2 is the same, then how could the utterance of D2 be sincere and that of D1 not—if, as A2 states, the sincerity of an assertion requires only that the asserter believe what she asserts? A view (perhaps like Perry’s (1979)), according to which my D1-belief and my D2-belief agree<sub>C</sub>, may be able to account for the differences between the two beliefs by pointing towards different ways or modes in which the shared content is believed. It may also be compatible with the supervenience of correctness on content, as expressed by hypothesis H. But it will have to abandon one of assumptions A2 and A3, or else bite the bullet and accept that whenever an utterance by MK of D2 is sincere, then so would have been an utterance of D1. On the assumption of framework A1–A4, and given that it seems clear enough that an utterance of D1 by me may be insincere when an utterance of D2 is not, I conclude that my D1-belief and my D2-belief have distinct contents, i.e. do not agree<sub>C</sub>.

Given the framework, then, the two beliefs have distinct contents. They also seem to agree<sub>N</sub>: it is necessary that if one is correct then so is the other. Now, if agreement<sub>N</sub> supervenes on content, how could the content of the D1-belief differ from the content of the D2-belief? I shall argue that within our framework, straightness (as expressed by hypothesis H) requires that the content of my D1-belief is a so-called “egocentric” content: a content that is only capable of being the content of the beliefs and assertions of one person, in this case me. In other words: we can preserve assumptions A1–A4 and the supervenience of correctness on content only at the cost of admitting egocentric contents. This can serve as a motivation for abandoning the supervenience hypothesis H.

Suppose for reduction that the content of the D1-belief is not egocentric, i.e. that it is capable of being believed by someone other than me. Then such a person could not express a belief with that content by uttering the same sentence, D1. For the belief another person would express by uttering D1 might be *incorrect* while my D1-belief is correct—they do not agree<sub>N</sub>. So, that belief by the other person could not share a content with my D1-belief, at least not if H is true. So, unless the other person’s belief in the content of my D1-belief is mysteriously inexpressible, there will be another sentence D3, by means of which the other person can express her belief in the content of my D1-belief. But whatever sentence could serve for that purpose, we will be able to devise a scenario *S\** just like *S*: It would be insincere for me to use D1, while it would not be insincere for me to use D3 (imagine, for example, that D3 is the sentence used by the announcer on the public address system). So, in the presence of H, contrary to the supposition, D3 cannot express the same content I express when I use D1. So the only way to maintain that my D1-belief and my D2-belief have distinct

---

the same content, and who holds on to (A2) would have to say either that it is not a lie, or that some lies are sincere.

<sup>14</sup> More precisely: the D1-belief together with a desire to avoid pain and certain typical means-ends-beliefs will be sufficient to motivate self-saving measures, while the D2-belief will not.

contents is to say that the content of my D1-belief is egocentric, i.e. is not available as a belief-content to anyone other than me.

Some theorists accept the existence of egocentric contents, for example Frege 1918, Evans 1981, Pollock 1982 or Stanley 2011. However, any such account must abandon the simple model of communication according to which communication centrally consists in asserting the contents one believes, and thereby making available these contents for the audience to believe. So, in so far as this simple picture provided a motivation for H, the egocentric account has to give up on it in any case. It is worth asking what else there remains to motivate H.

Other theorists, such as Egan 2007 and Moss 2011, give up two of the framework assumptions: A2 and A3. According to Egan, the content of the belief that would make my utterance of D1 sincere is not the same as the semantic content of the sentence D1, nor is it the same as the content of the assertion I thereby make. Thus Egan hopes to preserve part of the simple picture of communication, namely the part that says that communication centrally consists in asserting contents and thereby making them available for others to believe.

We can maintain the framework assumptions and avoid egocentric contents by abandoning H: if correctness does not supervene on content, then we can say that the content of the D1-belief is a public content, but one which it is not equally correct or incorrect to believe for everyone at every time. Rather, it is correctly believable only by those whose pants are on fire. When you believe this content, then your belief is correct if your pants are on fire, while if I believe it, then my belief is correct only if my pants are on fire.

This type of view preserves the idea that communication centrally consists in asserting the contents one believes, and thereby making available these contents for the audience to believe. Of course, the recommendation for the audience is not, as maintained by Egan, that one should come to believe the contents of those assertions that are by trustworthy sources. Rather, the recommendation should be more differentiated: that one should believe the contents asserted by trustworthy sources only when the source is relevantly similar to oneself. I will explain in the last section what “relevantly similar” needs to mean here.

There is another way of making the argument that *de se* contents require us to give up straightness.<sup>15</sup> Suppose Gustav Lauben is in Stanford Library, and he has access to all the location-neutral facts, i.e. the sort of facts that can be found in libraries and in Wikipedia. In other words, he has access to all the true contents the correctness of believing which does not depend on who believes it where and when. In other words, Lauben is fully informed about those contents that it is correct to believe, and which it is correct to believe no matter what the believer’s location is. Still, it would seem compatible with him having all that information that he does not have the belief he might express by saying

GL I am Gustav Lauben and I am now in Stanford Library.

Clearly, he knows a lot about Gustav Lauben, and his whereabouts at various times. But the content of the belief he would express by saying GL does not seem to be a content that it would be correct for anyone else to believe, or for Gustav Lauben at a different location. So the correctness of some beliefs does not seem to supervene on

---

<sup>15</sup> See Lewis 1979.

content, e.g. Gustav Lauben's GL-belief, the correctness of which does not supervene on content: if anyone else believed the very same content, they would be wrong.

The only way to avoid this conclusion would seem to be to insist that the content of Lauben's GL-belief is, after all, one that (per impossibile) it would be correct to believe for anyone, no matter who they are and where they are located. But in the presence of our framework assumptions and straightness, the only way to ensure that is to insist that no-one but Gustav Lauben at that time and place is able to think that content. Thus, we can again draw the conclusion that we can save straightness only at the cost of admitting egocentric contents.

#### **4. Limits of Straightness: Epistemic Modals, Taste, etc.**

There is another class of cases that undermine straightness. They involve the sorts of expressions that have recently been debated within the relativism-contextualism literature. I shall here focus just on one sample case, one involving epistemic modals. However, similar cases could be constructed using probability assigning sentences, those containing predicates of personal taste, or—more controversially—knowledge ascribing sentences or sentences expressing moral evaluations.<sup>16</sup>

Consider a situation *S* where Alfred's evidence does not rule out that Dan is in Paris, while his interlocutor Bert knows that Dan is in Barcelona (and that being in Barcelona rules out being in Paris). Now consider two potential utterances by Alfred in that situation, of each of the following two sentences:

E1     Dan might be in Paris.

E2     My evidence does not rule out that Dan is in Paris.

Given what Alfred knows, the assertions as well as the beliefs they express seem to be correct.<sup>17</sup> Now consider also Bert's potential responses to Alfred's utterances:

E1r    [responding to Alfred's E1] That's wrong.

E2r    [responding to Alfred's E2] That's right.

Bert's two responses are also correct: in *S* it is correct for Bert to respond to Alfred's utterance of E1 by saying, e.g. "That's wrong.", or perhaps "I disagree.". However, such responses would not be correct responses to Alfred's utterance of E2. For Alfred's evidence does not rule out Dan's being in Paris. Bert's response "That's right." is therefore correct, and so would "I agree.".

---

<sup>16</sup> For an overview of the issues and some references see Kölbel 2008 and forthcoming. Sample relativist accounts of epistemic modals are Egan 2007 and Mac Farlane 2011. Sample contextualist accounts are those given in von Fintel and Gillies 2008 and in Dowell 2011.

<sup>17</sup> Contextualists might deny that the belief expressed, and the assertion effected, by the utterance of E1 in *S* is correct. The idea is that a sufficiently sophisticated version of contextualism may predict that the contextually relevant knowledge includes Bert's knowledge, so that the content expressed by E1 in *S* is, roughly, a content that is true iff the knowledge shared between Bert and Alfred is compatible with Dan's being in Paris. In response, I would insist that the utterance of E1 in *S* at least *appears* to be correct, and that saying it is not correct is a theoretical cost, a cost that may well be compensated for by other advantages, and that may well be mitigated by some pragmatic explanation for the appearance of correctness. Another move would be to modify the example and consider assessments by eavesdroppers instead of Bert's E1r and E2r (see Egan 2007). In any case, the quick argument given here is merely meant to give a taste of the kind of motivation that motivates a relativist semantics for epistemic modals.

I want to argue that straightness causes difficulties for our understanding of Bert's responses E1r and E2r. What is it that Bert calls right and wrong respectively in his responses? What is it that "that" in E1r and E2r refers to? One simple and plausible proposal would be to say that "that" refers respectively to the contents of each of Alfred's utterances. Given the correctness of both of Bert's responses, this would mean that Alfred's utterances can't have the same content (the same content can't correctly be called right and wrong by the same person.) So let us first explore the option that "that" refers to the contents of Alfred's utterances, and that these contents are distinct.

Alfred's E1-belief is correct, and so is Alfred's E1-assertion. So, given straightness, if anyone else believed or asserted the content of Alfred's E1-belief (= the content of his E1-assertion), then they would also be doing so correctly. So, in particular, if Bert believed or asserted this content, he would do so correctly. But in that case, how can it be that Bert's response E1r is correct? If "that" refers to the content of Alfred's E1-belief, and it is correct to believe this content for everyone, including Bert, then what is the property of wrongness that Bert is correctly attributing to the content of Alfred's E1-belief? It would need to be a property that is compatible with correctness (as in conformity with the norms of correct application of concepts). It would need to be coherent for Bert to believe and assert a content and at the same time call that content "wrong". Calling a content "wrong" would not commit one to not believing that content. But that doesn't seem to be the correct interpretation of E1r: to say "That's wrong." commits one to not accepting whatever "that" refers to. I conclude that given straightness, "that" in Bert's utterance of E1r cannot refer to the content of Alfred's E1-assertion.

But what could "that" refer to, if not the respective contents of Alfred's utterances of E1 and E2? Given straightness, and given that Alfred believes the content in question correctly, "that" can't refer to the content Alfred believes. But could it perhaps refer to the content that Bert *would* express were he to utter E1? This would certainly square with the fact that Bert's response E1r is correct, for it is indeed wrong for Bert to believe the content he would express if he uttered E1. However, if the same construal of "that" were applied to E2r, we would get the wrong result: Bert should equally reject the content he would express were he to use E2.<sup>18</sup>

It seems difficult to account for the correctness of Bert's responses if we maintain straightness. However, dropping straightness seems to make everything fall neatly into place.<sup>19</sup> Suppose that Alfred's E1-belief has a non-portable content: a content that it is correct to believe for some people at some times in some places, but not by others or at other times or in other places. Then the content in question might simply be one that it is correct for Alfred to believe (given his state of knowledge before Bert's response), but that it is not correct for Bert to believe (given his state of knowledge). So Bert can reject the content of Alfred's E1-belief correctly, while Alfred accepts it correctly.

---

<sup>18</sup> Contextualists may be able to overcome this difficulty, so this is not a conclusive objection, and not meant to be one. For one thing, contextualists might point out that the interpretation of phrases like "That's wrong." is in any case quite complicated. It depends not only on the previous utterance and the features of context that I defined for *S* but rather on what is the question at issue. Thus, for example, if Alfred's utterance of E2 were to be a response to the question "Where is Dan?", then E2R would suddenly seem incorrect. Thanks to Craige Roberts for discussion.

<sup>19</sup> At least if we leave problems concerning the *de se* aside for a moment. More about this in a moment.

Suppose we say, by contrast, that the content of Alfred's E2-belief is fully portable: i.e. that if it is correct for one person to believe it at one time, then it is correct for anyone to believe it at any time. We might describe this content as the content that Alfred's evidence does not rule out Dan's being in Paris. Then we can explain the fact that Bert's contrasting response E2r is also correct.

However, the attentive reader will have noticed the snag: in the last section I argued that we should treat the content expressed by D1 ("My pants are on fire.") as a non-portable *de se* content, in order to avoid saying that it's egocentric, and as distinct from the content of my D2-belief ("MK's pants are on fire."), so that we could explain their different sincerity conditions. However, the same should apply to a sentence like E2: the content Alfred asserts with E2 stands to the content he asserts with

E3     Alfred's evidence does not rule out that Dan is in Paris.

as D1 stands to D2. If we are to go along with the proposal in section 3, we ought to treat E2 as expressing a non-portable *de se* content, and take this as what accounts for the difference from the content expressed by sentences like E3. But then the prediction seems to be that Bert's response E2r can't be correct, for if Bert is endorsing a *de se* content expressed by Alfred's utterance of E2, then it would seem that that's a *de se* content that it is correct to believe only for believers whose evidence does not rule out that Dan is in Paris—so not a content that Bert can correctly endorse.

Thus we have two separate reasons to give up straightness, but the corresponding solutions are not compatible with one another. Abandoning the doctrine that agreement supervenes on content does not by itself solve both problems. One way out is to maintain after all (as Perry 1979 does) that my D1-belief and my D2-belief agree in content (and that the differences between them are not differences in content). As we saw in section 3, this involves giving up framework assumptions A2 and A3. In the last section, I shall schematically outline a view according to which both E1 and D1 express non-portable contents. As we shall see, this will also involve abandoning A2 and A3 in some sense.

## **5. Living without Straightness**

I said at the beginning that straightness is a *prima facie* desirable feature of a theory of content: straight contents of belief and assertion are maximally "portable". Beliefs with such contents don't have a sell-by date. They do not require routine updating (updating in the absence of new evidence). The correctness of believing a portable content never expires due to a change in who believes it when, where and in what circumstance. Portability means that communication can consist simply in the broadcasting and receiving of contents, and information storage merely involves retaining these contents. But considerations of *de se* and epistemic modal beliefs suggest that maximum portability is not available. *De se* beliefs force us to say either that some contents are egocentric (the opposite of portable), or that many of the contents of our beliefs cannot be asserted: what we assert is then only a portable counterpart of what we believe, so that we reject assumption A2 (see especially Moss 2011). Giving up straightness means giving up on the general portability of contents and therefore on the idea of communication as the simple broadcasting and receiving of contents. But so do the alternative views.

However, how should we conceptualize communication, specifically assertoric communication, on a view that operates with contents that are not generally portable, so that the correctness of a belief or assertion is not determined by its content alone?

If we admit non-egocentric contents that are non-portable, i.e. contents that it is correct to believe for one thinker at one time yet not by another, or not at another time, then we think that the correctness of beliefs does not supervene on belief-content. In addition to depending on the content, the correctness of a belief can depend on what we can generally call the “location” of the believer. Now, I do not want to offer a theory here of what a location should be (see Liao 2012 for some general considerations on what a location might be). But for simplicity, let us assume that locations are fully determined by a thinker and a time, so let us think of locations as ordered pairs of a thinker and a time.

I want to propose two general models of how communication *could* proceed involving such contents: the *local portability model* and the *surrogate content model*. I shall be proposing that communication with “pure” indexicals follows the surrogate content model, while communication with non-portable contents of the type we encountered in section 4 above follows the local portability model.

*The local portability model:* Even if a content is not portable, it may well be *portable relative to a class of locations*:

- (R) A content  $c$  is portable relative to a class of locations  $L$  just if any belief with content  $c$  held at a location that is a member of  $L$ , agrees with any other such belief.

For example, the content of the belief I would express by uttering E1 is not generally portable, for it will be correct to believe that content for any believer whose evidence does not rule out that Dan is in Paris, but not correct to believe this for believers whose evidence does rule it out. The content is correct at some locations but not at others. However, the content is portable relative to any class of locations whose agent components have the same (or relevantly similar) evidence.

Now, the content of an assertion may be *locally portable* in the sense that it is portable relative to a class of locations of which the locations of the conversational participants are (or are presupposed to be<sup>20</sup>) members. If the local portability of a content is assumed or known by the participants of a conversation, then assertoric communication can proceed along the lines of the simple model (which we were hoping to conserve by maintaining straightness): the content believed by a speaker is asserted by her, and thereby made available for her audience to believe.<sup>21</sup>

Some non-portable contents allow little room for local portability. Thus, the content of my D1-belief will be portable relative to all locations  $\langle a, t \rangle$ , such that  $a = MK$  and  $t =$  the time of my D1-belief. But this means that it will not be locally portable except when all the participants are (or are presupposed to be) me. Other

---

<sup>20</sup> The brackets are meant to offer a non-equivalent variant, which may be preferable given the kind of model of conversation offered in Kölbel 2010 and 2011 (inspired by Stalnaker 1978 and 2002).

<sup>21</sup> This still leaves the question of how communications proceeds when contents are not locally portable. In the special case of epistemic modals, there may be an imperative for participants of a conversation mutually to adjust their evidence in such a way that epistemic modal contents are locally portable. In other cases (e.g. Lasersohn 2005’s “exocentric uses”) the presupposed location of participants may converge, even if actual locations diverge. But there are presumably also assertions that are unsuccessful in some sense because they are assertions of contents that are not locally portable.

non-portable contents are guaranteed to be locally portable, such as, for example, the belief anyone would express by “We are lost.”, at least on a reading of “we” according to which it picks out the participants of the conversation.

*The surrogate content model:* Our framework assumptions involved the idea that when an assertion is made, the content of the assertion is the semantic content of the sentence used at the context relevant to the assertion (assumption A3), and that an assertion is sincere just if the asserter believes what has been asserted (A2). On the surrogate content model, the contents that are broadcast and received in communication are not in general the contents believing which makes the relevant assertions sincere. Instead, when a non-portable content is asserted, the content communicated is a portable surrogate of the asserted content. For example, when I utter D1 (“My pants are on fire.”) sincerely, I believe and assert a non-portable *de se* content. This belief and assertion do not agree with my audience’s potential beliefs or assertions with the same content. However, there are surrogate portable contents that can play the role of the content that gets broadcast by me and received by my audience, and which in some sense encapsulates that part of my information that is independent of my location (or independent of those aspects of my location that are not shared by my audience).

A simple suggestion of a content that can play this role would be the content expressed by D2: the content one might describe as “the content that MK’s pants are on fire”. That content is portable (even fully portable, if we ignore questions of time). If all the participants of the conversation assume that I am MK, then this content can serve as the portable surrogate of D1’s non-portable content.

Sarah Moss (forthcoming) claims that there is a portable surrogate content  $q$  for every *de se* (non-portable) content  $p$  in the following sense:  $q$  is equivalent to  $p$  “given what I believe with certainty”. She tries to lend plausibility to this claim by pointing out that I can always stipulatively introduce a name “Dr Me” for myself, so that I believe with certainty that I am Dr Me. Thus any non-portable content expressible by a sentence ‘I ... .’ is equivalent, given what I believe with certainty, with the content expressible by ‘Dr Me ... .’. While this general approach makes some progress towards a solution to the problem of *de se* communication, something central is missing: it does not explain how any surrogate content generated in this way is going to figure in my communication with others who may not be aware of the stipulation I have made, the stipulation I would express by saying “I am Dr Me.”. Thus, even if the newly generated contents are equivalent, given what *I* believe with certainty, with the original *de se* content, they will not be equivalent given what my audience believes, unless they have similar access to the imagined stipulation. Suitable surrogate contents do not need to be equivalent *given what the speaker believes* with their originals, but rather, they need to be equivalent *given the information (assumed to be) available to all conversational participants*.

Suitable surrogate contents are available nevertheless. This much is guaranteed by general constraints on a functioning conversation. For participants of a conversation will always be able to identify the speakers of the various utterances of the conversation *in some way that is shared by all participants*. Perhaps in limiting cases, they know extremely little about who is speaking when they witness an utterance of “My pants are on fire.”. However, they can always think of the utterer of that utterance as “the person who just uttered ‘My pants are on fire.’”. Or perhaps, a little later, they can think of that person as “the person who a little while ago uttered ‘My

pants are on fire.”. For a conversation to be minimally functional, participants need to be able to keep track of what has been said and by whom—“by whom” in the minimal sense of being able to identify and re-identify the speakers of various utterances. When this minimal condition is not met, the conversation is dysfunctional and transmission of information via sentences that express non-portable contents is severely impaired.

For ease of exposition, let us pretend that the minimal way in which participants of a conversation can *always* think about their conversational partners is as “utterer of the first utterance of this conversation”, “utterer of the second utterance of this conversation”, etc. I.e. we are pretending that conversational participants can at least identify and remember the various utterances of the conversation by numbering them. If this is so, then there are indeed always suitable surrogate contents. Thus, if my utterance of “My pants are on fire.” is the seventh utterance of the conversation, then a locally portable content conveyed by the seventh utterance would be the content that would be expressed by “The pants of the utterer of the seventh utterance of this conversation are on fire.” (i.e. the content that would be expressed by the sentence were it uttered by any of the participants *of that conversation*). Now, this surrogate is not *fully* portable. It is a non-portable *de se* content. However, it is *locally portable*, and this is all we needed to explain how conversation proceeds.

On this view, even if the content of the belief that makes a certain assertion sincere is non-portable and not even locally portable, communication proceeds via a certain surrogate content, which, while being portable (or locally portable) nevertheless yields beliefs that agree with the belief that made (or would have made) the assertion sincere.

The way I sketched the surrogate content model, it did not involve questioning assumptions A2 and A3. Thus, for all I said, the content asserted could be identical to the semantic content of the sentence used for the relevant context, and moreover, it could be identical to the content of the belief which would make the assertion sincere. However, the content communicated may differ from this, for it may be a suitable surrogate content. This seems to follow the model of Stalnaker (1978): the diverging content (in Stalnaker’s case: the diagonal) is what updates the conversational context, but that does not mean it is the assertoric content.

A different terminological approach would seem equally reasonable: call the content that is used to update the conversational context the “assertoric content”. In this case, the surrogate content model involves saying that the surrogate content is the assertoric content, which is then distinct from the content of the belief that makes the utterance sincere. This involves rejecting A2 and A3, and seems to be the model proposed by Egan 2007.

How, then, should we construe assertoric communication if the cases of both section 3 and section 4 have moved us to accept a non-straight framework involving non-portable contents? The programmatic suggestion I would like to make is this: both *de se* beliefs and assertions and beliefs and assertions about what is epistemically possible should be construed as having non-portable contents. However assertoric communication follows the surrogate content model when sentences such as D1 or E2 are used. But it follows the local portability model when sentences like E1 are used. In other words, a sentence like E1 is, by virtue of its meaning, suited for conversations in which a convergence of the participants with respect to relevant evidence is presupposed. Sentences like D1 and E2, by contrast, are, by virtue of their

meaning suited to conversations where no relevant convergence of the locations of participants needs to be presupposed.

The proposal, thus, is that it is a feature of the conventional meaning of certain sentences that express non-portable contents, that they are suited for the surrogate content model of communication. Thus, what is seen to be the contribution to the conversational context is the surrogate content. It is the surrogate content which is treated as the content that has been put on the conversational table. By contrast, other sentences, such as E1, are suited (as a matter of meaning) for putting on the conversational table a non-portable content. This contrast finally explains the correctness of both Bert's evaluations E1r and E2r: in both cases, "that" refers to the content that has been put on the conversational table. In the case of E1r this is a non-portable content. But in the case of E2r, this is a portable surrogate.<sup>2223</sup>

### **References:**

- Cappelen, Herman and John Hawthorne (2009) *Relativism and Monadic Truth*, Oxford: Oxford University Press.
- Carnap, Rudolf (1955) Meaning and Synonymy in Natural Languages. *Philosophical Studies* 6, pp. 33–47.
- Dowell, J.L. 2011. "A Flexibly Contextualist Account of Epistemic Modals." *Philosopher's Imprint* 11, no. 14.
- Dummett, Michael (1991) *The Logical Basis of Metaphysics*. Cambridge, MA: Harvard University Press.
- Egan, A., J. Hawthorne and B. Weatherson (2005) Epistemic Modals in Context. In G. Preyer and G. Peter (eds), *Contextualism in Philosophy*, Oxford: OUP.
- Egan, Andy (2007) Epistemic Modals, Relativism and Assertion. *Philosophical Studies* 133, 1–22.
- Evans, Gareth (1981) Understanding Demonstratives. In H. Parret and J. Bouveresse (eds.), *Meaning and Understanding*, Berlin: W. de Gruyter. Reprinted in G. Evans, *Collected Papers*, Oxford: Clarendon Press, 291–321.
- von Fintel, Kai, and Anthony S. Gillies (2008) "CIA leaks". *Philosophical Review* 117, pp. 77–98.
- Frege, Gottlob (1892) "Über Sinn und Bedeutung". *Zeitschrift für Philosophie und philosophische Kritik* 100, pp. 25–50. English translation (partial) in Darragh

---

<sup>22</sup> This proposal offers a new answer to the problem of communication with "centered contents", as it has recently been discussed by, for example, Torre 2010, Ninan 2010a and 2010b, Stojanovic 2012, Kindermann forthcoming, or Weber forthcoming. It is novel in offering two distinct communicative models for two distinct types of centered content (classic de se contents and relativistic contents). It is also novel in showing how a surrogate content approach need not introduce fully portable surrogates. Rather, locally portable surrogates are sufficient.

<sup>23</sup> This paper has benefitted from comments and discussions at the workshop on *Disagreement* at the University of Tartu, Estonia, organized by Daniel Cohnitz, at the SIFA-Workshop on *Disagreement* in Turin, Italy, in December 2011 and the workshop on *Relativism and Disagreement* held at the Lichtenberg Kolleg, Georg-August-Universität Göttingen, Germany. The research leading to the results here presented was supported by the Spanish Ministry of Economy and Competitiveness: I+D Programme, projects FFI2009-13436 and FFI 2012-37658, Consolider-Ingenio Programme, Project CSD2009-0056, as well as the FP7 Programme of the European Union, project no. 238128.

- Byrne and Max Kölbel (eds), *Arguing about Language*, London: Routledge 2010.
- Frege, Gottlob (1918) *Der Gedanke. Eine logische Untersuchung. Beiträge zur Philosophie des deutschen Idealismus* 1, 58–77. English translation in Brian McGuinness (ed.), *Frege: Collected Papers on Mathematics, Logic and Philosophy*, Oxford: Basil Blackwell 1984.
- Kaplan, D. (1977) On Demonstratives. in J. Almog, J. Perry and H. Wettstein (eds), *Themes from Kaplan* (1989), Oxford: Oxford University Press, 481–563.
- Kindermann, Dirk (2012) “Varieties of Centering and the Common Ground”. Conference presentation at the workshop *Centered Content and Communication*, June 2012 in Barcelona.
- King, Jeffrey (2003) Tense, Modality, and Semantic Values. *Philosophical Perspectives* 17, 195–245.
- Kölbel, Max (2008) Motivations for Relativism. In Manuel Garcia-Carpintero and Max Kölbel (eds.), *Relative Truth*, Oxford: Oxford University Press, 1–38.
- Kölbel, Max (2009) “The Evidence for Relativism”. *Synthese* 166, pp. 375–95.
- Kölbel, Max (2010) “Vagueness as Semantic”. In R. Dietz & S. Moruzzi (eds), *Cuts and Clouds: Issues in the Philosophy of Vagueness*, Oxford: Oxford University Press, pp. 304–26.
- Kölbel, Max (2012) “Should we be Pluralists about Truth?”. In Nicolaj Jang Pedersen and Cory Wright (eds), *Truth Pluralism: Current Debates*, Oxford: Oxford University Press 2012.
- Kölbel, Max (forthcoming) “Relativism”. Forthcoming in *Philosophy Compass*.
- Künne, Wolfgang (1992) “Hybrid Proper Names”. *Mind* 101, pp. 721–31.
- Laserson, Peter (2005) Context Dependence, Disagreement, and Predicates of Personal Taste. *Linguistics and Philosophy* 28, 643–86.
- Lewis, David (1979) “Attitudes De Dicto and De Se”. *Philosophical Review* 88, 513–543.
- Lewis, David (1986) *On the Plurality of Worlds*. Oxford: Blackwell.
- Liao, Shen-Li (2012) What are Centered Worlds. Forthcoming in *Philosophical Quarterly*.
- MacFarlane, John (2007) Relativism and Disagreement. *Philosophical Studies* 132, 17–31.
- MacFarlane, John (2009) Epistemic Modals are Assessment-Sensitive. In Brian Weatherson and Andy Egan (eds), *Epistemic Modals*, Oxford: Oxford University Press.
- Moss, Sarah (forthcoming) Updating as Communication. *Philosophy and Phenomenological Research*.
- Nelson, Michael (2002) “Puzzling pairs”. *Philosophical Studies* 108, 109–19.
- Ninan, D. (2010a) Semantics and the Objects of Assertion. *Linguistics and Philosophy* 33, 355–380.

- Ninan, D. (2010b) De Se Attitudes: Ascription and Communication. *Philosophy Compass* 5, 551–67.
- Perry, John (1977) “Frege on demonstratives”. *Philosophical Review* 86, pp. 474–97.
- Perry, John (1979) “The Problem of the Essential Indexical”. *Nous* 13, pp. 3–21.
- Pollock, J. (1982) *Language and Thought*. Princeton University Press
- Rabern, Brian (forthcoming a) “Against the Identification of Assertoric Content with Compositional Value”, *Synthese*.
- Rabern, Brian (forthcoming b) “Propositions and Multiple Indexing”, *Thought*.
- Recanati, François (2007) *Perspectival Thought*, Oxford: Oxford University Press.
- Salmon, Nathan (1986) *Frege’s Puzzle*. Cambridge, Massachusetts: MIT-Press.
- Schaffer, Jonathan (forthcoming) “Necessitarianism”. *Synthese*.
- Soames, Scott (2002) *Beyond Rigidity*. Oxford: Oxford University Press.
- Stalnaker, Robert (1978) Assertion. In P. Cole (ed.), *Syntax and Semantics 9: Pragmatics*, 315–22.
- Stalnaker, Robert (1984) *Inquiry*. Cambridge, MA: MIT-Press.
- Stalnaker, Robert (2002) Common Ground. *Linguistics and Philosophy* 25, 701–721.
- Stalnaker, Robert (2008). *Our Knowledge of the Internal World*. Oxford University Press.
- Stanley, Jason (1997) “Names and Rigid Designation”. In Bob Hale and Crispin Wright (eds.), *A Companion to the Philosophy of Language*. Oxford: Blackwell, pp. 555–85.
- Stanley, Jason (2000) “Context and Logical Form”. *Linguistics & Philosophy* 23, pp. 391–434.
- Stanley, Jason (2007) “Introduction”. In Jason Stanley, *Language in Context: Selected Essays*. Oxford: Oxford University Press.
- Stanley, J. (2011) *Know How*. Oxford: Oxford University Press.
- Stojanovic, Isidora (2007) “Talking about Taste: Disagreements, Implicit Arguments and Relative Truth”. *Linguistics and Philosophy* 30, pp. 691–706.
- Stojanovic, Isidora (2012) “The Problem of De Se Assertion”. *Erkenntnis* 76, pp. 49–58.
- Sundell, Tim (2012) “Disagreements about Taste”. *Philosophical Studies* 155, pp. 267–88.
- Torre, Stephan (2010) Centered Assertion. *Philosophical Studies* 150, 97–114.
- Weber, Clas (2012) “Centered Communication”. Conference presentation at the workshop *Centered Content and Communication*, June 2012 in Barcelona. Available at <http://www.clasweber.net/>.