

Centered Communication

Clas Weber
School of Philosophy
Australia National University
Clas.Weber@anu.edu.au

According to an attractive account of belief, our beliefs have centered content. According to an attractive account of communication, we utter sentences to express our beliefs and share them with each other. However, the two accounts are in conflict. In this paper I explore the consequences of holding on to the claim that beliefs have centered content. If we do in fact express the centered content of our beliefs, the content of the belief the hearer acquires cannot in general be identical to the content the speaker expresses. I sketch an alternative account of communication, the *Recentering model*, that accepts this consequence and explains how expressed and acquired content are related.

1 Introduction

There is good reason to think that some of our beliefs have *centered* content. It is furthermore plausible to assume that we can express these beliefs using ordinary sentences of English. But there is a problem: the assumption that we can express centered beliefs clashes with a natural account of communication. According to this account, the speaker is able to directly pass on the content of her beliefs to the hearer. This picture of communication cannot be upheld in the face of centered content. Assume, for instance, that my utterance of “I have Groat’s disease” expresses the content of my centered belief that *I have Groat’s disease*. The account of communication predicts that if you accepted my assertion and came to endorse its centered content, you would acquire the belief that *you yourself* have Groat’s disease. This is obviously the wrong result. What you in fact come to learn is that *I*, the speaker, have Groat’s disease. Hence, if we are indeed able to express our centered beliefs, the picture of communication has to be modified. The content you acquire from my utterance cannot be the content I expressed.

In this paper I sketch an alternative account of communication, the *Recentering model*, that endorses this consequence and explains how the expressed content and the acquired content are related. According to the *Recentering model*, it is still the case that the speaker literally expresses one of her beliefs. The belief the hearer acquires, however, is not the one expressed

by the utterance.¹ Rather, she acquires a content that is related in a certain way to the utterance content. The acquired content is determined by the content expressed by the speaker together with the hearer's beliefs about how she is related to the speaker's context.

First, I introduce the centered account of belief and the standard model of communication and explain how the conflict between the two arises. The main objective of this paper is to show that we can hold on to the centered belief account, as there is an alternative model of communication: the Recentering model. Finally, I present an attempt to reconcile centered beliefs with the standard model of communication.

2 Centered Belief

Some of our beliefs concern what the universe is like. According to my belief that electrons exist, our universe contains electrons somewhere or other. Yet, not all of our beliefs are about the way the world is—many concern our place within it. They are about how things are with ourselves, who we are, and where we are located in space or time. If a belief concerns the world as a whole, we can explain what it is for someone to have such a belief in terms of her bearing a certain relation to a set of possible worlds (or a structured proposition).² John Perry (1979) and David Lewis (1979) have pointed out that this explanation fails for the second, self-locating kind of belief. Self-locating beliefs cannot be fully understood as binary relations between subjects and standard propositions.

The case against the standard account of belief can be illustrated by Lewis' story of the two gods (Lewis, 1979). The two gods live on different mountains. Even though they know exactly which possible world they inhabit—they are omniscient with respect to possible worlds propositions—they still suffer a certain kind of ignorance. Neither god knows on which mountain he lives. Hence, there is a kind of belief that cannot be modeled as a relation to a standard proposition.³

While Perry left standard belief content untouched and added a further element to the framework, so-called *belief states*, Lewis suggested to give up the standard account of content. Rather than identifying contents with sets of possible worlds, we should think of them as sets of possibilities of any kind (i.e. Lewisian properties), or equivalently as sets of *centered worlds* (world, time, individual triples). My belief that *I have Groat's disease* can be represented by the set of time-slices of world-bound individuals with Groat's disease, or equivalently as the set of centered worlds in which the individual at the center has Groat's disease at the time of the center. Within this framework, belief can be comprehensively explained in terms of binary relations between subjects and *centered propositions*. The attraction of Lewis' account is obvious: we

¹At least this is true for the cases we are interested in, i.e. the expression of indexical beliefs. For non-indexical beliefs everything can remain just as before.

²Modelling contents as sets of possible worlds isn't satisfactory for all theoretical purposes. The reason is the familiar problem of logical omniscience: logically equivalent contents are identified. This problem isn't solved by moving to centered content. I will nevertheless put this issue to the side.


³Stalnaker (2008) denies that one can be omniscient with respect to possible worlds propositions and at the same time suffer from self-locating ignorance. However, his response seems to rely on a doxastic form of haecceitism.

have a simple and unified picture of belief, which is a generalization of the standard account.⁴

One important motivation for rejecting Lewis' centered account has been the assumption that it does not allow for a plausible explanation of communication (Stalnaker, 1981), (Stalnaker, 2008).⁵ This paper aims to undermine this motivation. By presenting a centered account of communication that complements the centered belief framework, it demonstrates that the proponent of this framework can avail herself of a viable account of communication as well.

3 The FedEx Model of Communication

Intuitively, the fundamental point in making assertions is to express one's beliefs and to share them with the audience. A good theory of communication should account for this. The most straightforward way of doing so is to identify the content of an utterance with the content of the belief the speaker is verbalizing and which the hearer acquires. As Stalnaker puts it: "The objects or contents of belief are also the objects or contents of speech acts [...]" (Stalnaker, 1988, p. 151). According to this picture, successful communication proceeds in the following way: the speaker expresses one of her beliefs by uttering a sentence. The hearer understands the sentence, trusts the speaker and then comes to believe the content of the utterance herself. Speaker and hearer end up sharing the same belief. As the content of the speaker's belief—a certain package of information—is sent via the utterance to the hearer, we can label this account of communication the *FedEx model of communication*.⁶ Here is an illustration of the model:



One may characterize the core of the FedEx model in the following way:

The FedEx model

Cases of successful communication obey the following two principles:

1. *The mind-to-speech principle*: the content of the utterance is identical to the content of a belief the speaker expresses.
2. *The speech-to-mind principle*: the content of the utterance is identical to the content of a belief the hearer acquires.

⁴There are certain technical problems with the framework interpreted as a model for all attitudes (Nolan, 2006), (Ninan, 2008). A note on terminology: even though strictly speaking all content is now conceived of as *centered* (as sets of possibilities), I will nevertheless continue to talk of *uncentered content*—by that I am referring to centered propositions which correspond to standard possible worlds propositions.

⁵A related objection to the framework has been made with respect to a theory of belief updating, e.g. (Stalnaker, 2008). For a centered model of belief updating see (Meacham, 2010), (Schwarz, 2012).

⁶This label is stolen from Sarah Moss (2010) who calls it "the package delivery model".

The *mind-to-speech principle* takes us from the content of one of the speaker's beliefs to the content of the utterance. The *speech-to-mind principle* takes us from the utterance content to the content of a belief the hearer acquires. In conjunction, the two principles take us from the mind of the speaker to the mind of the hearer. They imply that in cases of successful communication, speaker and hearer come to share a belief whose content is identical to the utterance content. We might call this the *content identity principle*.⁷

The FedEx model is an attractive picture of communication. It is neat, simple and captures our central intuitions about communication. A very elegant implementation of the model has been given by Stalnaker (1978). Unfortunately, we get into trouble if we combine the FedEx model with the centered belief account.

4 The Dilemma

We were offered an attractive account of belief, together with an appealing model of communication. Unfortunately, it seems that we cannot accept both at the same time. What exactly is the difficulty in combining the centered belief account and the FedEx model? The problem is that if we plug centered content into the FedEx model, the model makes false predictions about what are successful and what are unsuccessful cases of communication.⁸

To see the problem, consider the following examples. A believes that she has Groat's disease and tells B: "I have Groat's disease". If everything goes well, we expect B to learn that A (or the speaker) has Groat's disease, while remaining agnostic about whether she herself has the disease. Call this case *Good Groat's*. In the second somewhat bizarre scenario B acquires the belief that she herself has Groat's disease, while remaining agnostic about whether A is diseased. Call this situation *Bad Groat's*. Clearly, communication has failed in *Bad Groat's*. However, a centered FedEx model would classify the cases in the opposite way. According to the centered belief account, the content of A's belief is the set of individuals with Groat's disease. On the assumption that she can express her belief with the sentence "I have Groat's disease", that is also the content of her utterance. As B shares a belief with the same content in *Bad Groat's* but not in *Good Groat's*, the FedEx model is obeyed in *Bad Groat's* but not in *Good Groat's*. This means that the former case gets misclassified as successful and the latter as unsuccessful communication. These examples aren't special cases—with the help of self-locating beliefs, we can produce problem cases *ad libitum*. Something has to give. As it isn't a viable option to simply reverse our classification, we seem forced to give up either the centered belief account or the FedEx model.

Strictly speaking, it isn't the combination of the centered belief account and the FedEx model by itself that is incompatible with our assessment of the cases. We have assumed that in the above

⁷The two principles aren't strictly necessary or sufficient for successful communication. For instance, there are cases of successful communication, in which the hearer merely recognizes what the speaker wants her to believe, without actually coming to endorse it. Moreover, it seems that even when she does, she has to do so *because* of the speaker's assertion, and not just by luck. These complications will not concern us here.

⁸That there is a problem in combining the two accounts was first pointed out by Stalnaker (1981). Egan (2007) provides a nice presentation of the problem as it arises within Stalnaker's implementation of the FedEx model.

cases A can in fact express the content of her belief—i.e. that centered beliefs are expressible in English. Let us label this assumption the *Expressibility of Centered Content*. One might try to save both the FedEx model and the centered belief account by rejecting this assumption and restricting the class of beliefs we can express: only uncentered beliefs (or the uncentered parts of centered beliefs) are expressible.⁹ In §7, I discuss the prospects of this strategy. Let us for now assume that we are able to give voice to the centered content of our beliefs (in case they have any).

A common reaction to the conflict has been to abandon the centered account of belief, rather than to modify the model of communication. *Prima facie*, this is somewhat surprising as belief seems to be the more fundamental phenomenon. Here, I want to explore a different path by presenting an alternative account of communication that allows us to retain the centered belief account.¹⁰

5 The Recentering Model of Communication


5.1 Expressed Content and Acquired Content

The Recentering model holds on to the *mind-to-speech principle*, but abandons the *speech-to-mind principle*.¹¹ According to that suggestion, we can preserve a central intuition about communication: the speaker can straightforwardly express the content of her beliefs. The hearer, however, doesn't acquire that same content. Rather, she adopts a different, though related, content. As a result, the *content identity principle* fails and we end up with two contents: i.) the *expressed content*, which is the content of the utterance and corresponds to the content of the belief the speaker is expressing; and ii.) the *acquired content*, which is the content the hearer acquires by accepting the utterance. The basic idea of the Recentering model can be pictured in the following way:

⁹Andy Egan (2007) and Sarah Moss (2010) have suggested a response along these lines. More accurately, the relevant claim should be labelled the *Expressibility of Malignant Centered Content*. The above problem doesn't arise for each and every piece of centered information. There are cases of centered content where everything would go well. One might then distinguish between *benign* centered content, which is expressible, and *malignant* centered content, which isn't. This is in fact the approach of Egan (2007). As my discussion of this strategy is independent of this distinction, we can here skip over this subtlety.

¹⁰After completion of the manuscript, it has been pointed out to me that both Neil Feit (2008) and Irene Heim (2004) sketch a model like this, without endorsing it. After an interesting discussion Feit opts instead to identify utterance content with the uncentered semantic value of the uttered sentence. I think this proposal doesn't work for the reasons given in §7. Frank Jackson (2010) also provides an account of communication that involves centered utterance content. However, within his framework the centered worlds which model a sentence's content don't always represent something about the speaker; e.g. for sentences involving names, the center represents the referent of the name. As a consequence, both principles of the FedEx model fail within Jackson's account. For the relevant cases, neither speaker nor hearer have a belief that corresponds to the utterance content. Another proposal for handling indexical utterances was suggested both by Stephan Torre (2010) and Dilip Ninan (2010a). They argue that we should understand utterance content in terms of *sequenced worlds*, i.e. centered worlds which allow for infinite sequences of individuals at the center. This strategy is hinted at in (Stalnaker, 2008, pp. 73-74). As Ninan (2010b) proposes to also assign sequenced content to beliefs, he has the prospect of maintaining the FedEx model. The cost, however, is that Ninan endows every subject with her own primitive belief relation.

¹¹As we will see in §6, it may be possible to do things the other way around.



The reason why we have to distinguish these two pieces of content lies in the notion of representation that underlies the centered framework. On the standard account of content, to make an assertion, or to have a belief, is to represent something about what the world is like. A hearer who trusts the speaker believes that her assertion is true, i.e. that the world is as the speaker represents it to be. Since hearer and speaker share the same world, she can directly take in the information the speaker expressed. In contrast, according to the centered account, beliefs represent something about the centered world of the thinker. Likewise, a centered assertion will convey information not just about the actual world, but also about the speaker's context.¹² Even though hearer and speaker share the same world, they don't share the same *centered* world. Therefore, the hearer, in trusting the speaker, cannot directly take the information expressed by the speaker to correctly characterize her own situation. Rather, trusting the speaker on the centered account is to believe that she said something which is true at *her* (i.e. the speaker's) context. Gaining information about one's own context from an utterance requires relating oneself to the utterance context.

This gives us an idea of how the hearer arrives at the acquired content from the expressed content. Two main factors determine the content the hearer can acquire from an utterance:

1. The content of the utterance, i.e. the expressed content.
2. The hearer's beliefs about her relation to the speaker (the utterance context).

It may seem as though we don't get a general, context-independent recipe for how to determine the acquired content associated with an utterance, since the acquired content depends partially on the hearer's assumptions about her relation to the speaker. And these assumptions may vary for the same hearer in different situations and also for different recipients of the same utterance token. That notwithstanding, it seems possible to reconstruct an intra- and interpersonally stable form of acquired content. It is important to note that the Recentering model doesn't demand of the hearer to locate the speaker in an *absolute* sense. The hearer doesn't have to possess information about the objective spatio-temporal location of the utterance context. Rather, she only has to locate the utterance context in relation to her own position. Arguably, there is a relation that the hearer can always fall back on. How little else she may know about the speaker,

¹²I am assuming that, metaphysically, contexts are centered worlds. Contextual parameters are determined by the features of the respective centered world (Lewis, 1980). Again, we can translate talk about *centered worlds* back into talk about Lewisian *possibilia*. The centered world that is the *utterance context* corresponds to *the speaker*, understood as the time-slice of the world-bound individual that produced the utterance. In the following I will go back and forth between the two formulations.

she can always locate the speaker with the following relation R : *h and s stand in R just in case s produced the utterance token that h is perceiving*. Not only does this give us a guarantee that the hearer always possesses the required locational information, we also get a common and stable form of acquired content, as every hearer can rely on this relation.¹³

How is the expressed content of an utterance related to the semantic value of the uttered sentence? The expressed content is the diagonal of the sentence's character (Kaplan, 1989). The diagonal of a character is a set of contexts. As we conceive of contexts as centered worlds, the diagonal corresponds to a set of centered worlds, i.e. a centered proposition. The diagonal of "I have Groat's disease" picks out the set of centered worlds in which the individual at the center has Groat's disease at the time and world of the center. As desired, this is the content of the corresponding 1st personal belief. Even though the expressed utterance content isn't identical to the semantic value of the corresponding sentence, it is conventionally associated with the utterance and can be grasped merely in virtue of one's linguistic competence.

5.2 The Model

The following is a general characterization of how the Recentering model works. The Recentering model differs from the FedEx model in an interesting way only in the case of centered content—for ordinary uncentered propositions, the two are equivalent.

The Recentering Model

1. The hearer perceives an utterance "u". [Perceiving]
2. The hearer believes that the expressed content of "u" is true of the speaker. [Centering]
3. The hearer believes that she is R -related to the speaker. [Locating]
4. The hearer infers information about herself from 2. & 3. [Recentering]

The process starts with the speaker producing an utterance and the hearer perceiving it. In the standard case, the hearer is in spatial and temporal proximity to the speaker and able to witness the speaker performing the utterance. Even when that isn't the case, she can always locate the utterance context using the above-mentioned relation R . Typically, she will learn more from an utterance, than just the fact that the person R -related to herself is ϕ . For instance, she may additionally know that the person R -related to her is Juri. Then she will learn that Juri is ϕ by learning that the person R -related to her is ϕ .

We can decompose the *Centering* step into two separate steps, which we might call *Understanding* and *Trusting*. Firstly, to gain information from an utterance, the hearer has to understand it. For us this means that she has to know what the expressed content of "u" is. Secondly, in the standard case of communication, the hearer will trust the speaker, i.e. she will believe that

¹³Admittedly, there are cases that cannot be handled that straightforwardly, e.g. utterance tokens that don't have a unique author, utterance tokens which are being re-used, i.e. for which the context of utterance doesn't coincide with the context, in which the utterance was produced etc. This suggests that things need to be refined. Still, I take it that these cases don't present a special problem for the Recentering model.

the speaker has made a true utterance.¹⁴ On the centered account, assertions represent something about the context in which they are produced. Then, for an utterance to be true, the utterance context has to be as the utterance represents it. To trust the speaker is to believe that the utterance correctly characterizes the speaker's context. Putting the *Understanding* and the *Trusting* step together, we get the *Centering* step: the hearer believes that the expressed content of “u” is true of the speaker's context.

We already noted that, since the speaker's and the hearer's centered worlds diverge, the hearer cannot simply take the expressed content to provide an accurate representation of her own context. Rather, she has to infer information about her location from information about a different location. For that, she has to locate the utterance context.

Finally, the hearer derives information about her own context from the information expressed by the speaker and her assumptions about her relationship to the utterance context. One may picture the entire process as the hearer going through two subsequent deductive arguments:

Centering

1. *Understanding premise*: “u” is true of the set of ϕ -individuals.

2. *Trusting premise*: “u” is true of the speaker.

Centering conclusion: The speaker is a ϕ -individual.

Recentering

1. *Centering premise*: The speaker is a ϕ -individual.

2. *Locating premise*: I am *R*-related to the speaker.

Recentering conclusion: I am *R*-related to a ϕ -individual.

We can say that the hearer first *centers* the utterance content on the speaker, and then *recenters* the inferred content on herself.

5.3 Examples

To get a better grip on the mechanics of the model, I will go through some examples. In the first example, I make explicit the otherwise contracted *Centering* step. The italicized parts represent the content of the beliefs of the hearer—as all beliefs are now conceived of as self-locating, they can all be paraphrased as: *I am such that ...* And since they are meant to capture the hearer's cognitive perspective, the description “the speaker” stands for something like: *the individual that produced the utterance which I am presently perceiving*. This way of referring to the speaker corresponds to employing relation *R* from §5.1. In case the hearer also employs relation *R* to locate the speaker in the *Locating* step, this step is more or less trivial and already implicit in the *Centering* step. However, that is not the case when more informative relations than *R* are being used to locate the utterance context. To make things more vivid in the following examples, I will illustrate the model with more substantive relations in the *Locating* step, rather than *R*.

In the first example, A tells B: “You have Groat's disease”. B rightly believes that she is the addressee of A's utterance. We expect her to acquire the 1st personal belief that she has Groat's

¹⁴As we noted in footnote 7, this isn't always the case. We might then want to replace the *Trusting* step with something weaker: e.g. *the speaker wants me to believe that she has made a true utterance*.

disease.

Example 1

Centering

1. B perceives the utterance “You have Groat’s disease” (being produced at the speaker’s present spatio-temporal location). [Perceiving]
2. B understands “You have Groat’s disease”: *the set of individuals such that “u” is true of the set of individuals whose addressee has Groat’s disease.* [Understanding]
3. B believes that “You have Groat’s disease” is true of the speaker: *the set of individuals such that “u” is true of the speaker.* [Trusting]
4. B believes that the expressed content of “You have Groat’s disease” is true of speaker: *the set of individuals such that the addressee of the speaker has Groat’s disease.* [Centering]

Recentering

1. B believes that the expressed content of “You have Groat’s disease” is true of the speaker: *the set of individuals such that the addressee of the speaker has Groat’s disease.* [Centering]
2. B believes that she is the addressee of the speaker: *the set of individuals such that they are the addressee of the speaker.* [Locating]
3. B infers information about herself from 1. & 2.: *the set of individuals with Groat’s disease.* [Recentering]

According to the Recentering model, B comes to self-locate amongst the individuals with Groat’s disease—she learns that she has Groat’s disease. This is the desired result.

In the second example, A asserts: “It is noon”. B directly perceives A performing the utterance and thereby knows that the utterance time is identical with the time when she encounters the utterance token. We expect B to learn that it is noon.

Example 2

1. B perceives the utterance “It is noon” (being produced at the speaker’s present spatio-temporal location). [Perceiving]
2. B believes that the expressed content is true of the speaker: *the set of individuals such that the speaker is a noon individual.* [Centering]
3. B believes that the utterance time is now: *the set of individuals that are simultaneous with the speaker.* [Locating]
4. B infers information about herself from 2. & 3.: *the set of noon individuals.* [Recentering]

The Recentering model predicts that B locates her present time-slice amongst the noon individuals, i.e. she comes to believe that it is noon. Again, the model delivers the right result.

For the last example, imagine that A and B are walking in the fog through a mine field. B doesn't know A's exact location. Still, as she can hear A's voice, she knows that A has to be nearby. A utters: "There is a mine over here". We expect B to learn that there is a mine nearby.

Example 3

1. B perceives the utterance "There is a mine over here". [Perceiving]
2. B believes that the expressed content is true of the speaker: *the set of individuals such that the speaker is co-located with a mine*. [Centering]
3. B believes that she is near the speaker: *the set of individuals that are near the speaker*. [Locating]
4. B infers information about herself from 2. & 3.: *the set of individuals that are near a mine*. [Recentering]

Even though B does not know the exact location of the speaker, she can still gain valuable information from her utterance. Moreover, the information that B acquires according to the model seems to be intuitively just the kind of information available to someone in B's situation.

6 Questions

There are some open questions concerning the Recentering model. I hope that my responses will help to further clarify the model.

1. *Question:* The Recentering model keeps the *mind-to-speech principle*, but abandons the *speech-to-mind principle*. Can we do things the other way around?

Response: Yes we can. In principle, we could identify utterance content with the content the hearer acquires, and thereby preserve the *speech-to-mind principle*. As the acquired content is often different from the content of the speaker's belief, we have to give up the *mind-to-speech principle*. The utterance "I have Groat's disease" will now be assigned the content the hearer acquires, i.e. the set of individuals such that the subject *R*-related to them has Groat's disease (it could be roughly paraphrased as *the person talking to me has Groat's disease*). Plausibly, relation *R* should be the one introduced in §5.2—otherwise, different hearers will sometimes employ different relations and we would have to ascribe numerous contents to the same utterance token corresponding to the distinct contents the different hearers acquire.

An advantage of the hearer-centric alternative may be that it better meshes with a Gricean metasemantics, according to which utterance content is determined by the information the speaker intends to instill in the hearer. A downside is that the connection between utterance content and semantic value of the uttered sentence is not as straightforward as on the Recentering model (there, utterance content was just the character's diagonal). More importantly, we would have to give up the central intuition that the speaker takes what she asserts to be true. To me, this intuition seems stronger than the one corresponding to the *speech-to-mind principle*, i.e. that the

hearer can directly take in the asserted information. The Recentring model therefore appears to be more natural than the hearer-centric alternative.

2. *Question:* The Recentring model seems rather complicated. Is it really plausible that we go through all the steps described and acquire information from an utterance by drawing certain inferences?

Response: Firstly, the model is not intended as a reconstruction of the *conscious* processes a hearer runs through when interpreting an utterance. It may very well be that most of it occurs at a sub-personal level.¹⁵

Secondly, in most cases, the derivation is straightforward. Typically, hearer and speaker share not only the same world, but also the same time parameter. Moreover, the fact that they do is transparent to them. For cases of temporal indexicality, the hearer can then take the utterance content as directly giving information about her own context. And even though speaker and hearer don't typically share exactly the same spatial location, they are usually in close proximity, such that the difference doesn't matter for most practical purposes.¹⁶ Hence, in many cases, the cognitive work for the hearer is rather minimal. In cases that aren't like these, on the other hand, it doesn't seem unrealistic that getting at the relevant information involves some cognitive work.

Thirdly, the uncentered account doesn't seem better off than the Recentring model. To see this, consider an utterance of "Juri is walking now". According to the uncentered account, its content is an eternal proposition (King, 2007). Which eternal proposition the utterance expresses depends on the time of utterance. Uttered at t_1 , the speaker expresses the proposition that *Juri is walking at t_1* . On this picture, the hearer has to infer what proposition is expressed from information about the time of utterance. Notice also that in contrast to the Recentring model, the uncentered model requires the hearer to locate the utterance context absolutely, and not just relative to her own position.¹⁷ Furthermore, on the Recentring model, the utterance content itself is *context-independent*—it is the diagonal of the sentence's character. The expressed content of "Juri is walking now" is the set of centered worlds at which Juri is walking at the time of the center, irrespective of when that sentence is uttered. The hearer can grasp the utterance content without possessing any information about the actual utterance context. These considerations suggest that the Recentring model doesn't fare worse than the uncentered account. In certain respects, the uncentered account imposes more challenging demands on the hearer.

3. *Question:* Someone may accept that the Recentring model accounts for the expression and acquisition of centered information, but still think that there is a difficulty: how can the account handle our judgments about disagreement, same-saying, etc.? Here, I will focus on the

¹⁵In fact, the hearer might not even have to go through anything like a subconscious inference to get at the acquired content from the information provided in step 2. and 3. I am sympathetic to a holistic picture of belief, according to which "beliefs" may be a "bogus plural" (Lewis, 1995, p. 423). On that picture, once the hearer has acquired the information of step 2. and 3., she will thereby have gained the acquired content as well, (unless that information is compartmentalized).

¹⁶For instance, imagine A and B riding together in a car, and A saying: "The movie theatre is just a couple of blocks ahead".

¹⁷The required information is quite often exactly the information the speaker is trying to convey, e.g. in the case of "It is t_1 now". The hearer can figure out what the speaker asserted only by antecedently knowing what time it is. But then, she cannot learn anything interesting from the utterance.

case of disagreement, assuming that the discussion carries over to the other notions.¹⁸ Consider the following case:

¬Disagree A utters: “I have Groat’s disease”. B responds: “I don’t have Groat’s disease.”

Intuitively, A and B don’t disagree. Contrast this with the following situation:

Disagree A utters: “I have Groat’s disease”. B responds: “You don’t have Groat’s disease.”

In this case, A and B do in fact disagree. The uncentered account can straightforwardly explain these judgements, as the expressed uncentered propositions are compatible in *¬Disagree*, while they are incompatible in *Disagree*. The centered account, on the other hand, seems to get things utterly wrong. The centered propositions expressed in *¬Disagree* are incompatible, while they are compatible for *Disagree*. Hence, the account wrongly predicts a disagreement in the first case and a lack of disagreement in the second.

Response: The first response on behalf of the Recentring model is to point out that the proponent of that model might be able to avail herself of uncentered content as well. Centered utterance content was obtained from the diagonal of the sentence’s character. Standard uncentered content is determined by evaluating the sentence’s character at the actual utterance context. We may accordingly associate sentences with both centered *and* uncentered content (Jackson, 1998), (Chalmers, 2004). Our judgments about disagreement can then track the uncentered content of the uttered sentence.¹⁹

This response may not go far enough, as the *mind-to-speech principle* (which the Recentring model wanted to retain) doesn’t seem to be upheld on a view where utterances have both centered and uncentered content, whereas beliefs are assigned only centered content. We then seem driven to also associate beliefs with both kinds of content, as e.g. Chalmers (2002) does. Here, I want to look at a different response. We might be able to give an account of disagreement purely in terms of centered content. The central idea behind the centered account of disagreement is to not focus directly on the content of utterances, but to rather compare the *contents of the beliefs the subjects associate with the utterances*. A subject associates a belief with an utterance, if, in case she is the producer of the utterance, it is the belief she expresses, or, in case she is the recipient of the utterance, it is the belief she would acquire in virtue of accepting the utterance. Consider once more *¬Disagree* and focus on the contents of the beliefs associated with the two utterances. The belief A associates with the first utterance is: *I have Groat’s disease*, and with the second she associates: *my addressee doesn’t have Groat’s disease*. It isn’t hard to see that the two beliefs are compatible. (The same goes for B, for whom we have: *my addressee has Groat’s disease* and *I don’t have Groat’s disease*). Hence, an account that explicates disagreement not

¹⁸Also, judgments about disagreement are more robust, than the highly flexible ones involving same-saying.

¹⁹I doubt that we will get a completely satisfactory account of disagreement by restricting ourselves to uncentered content. Here are some problematic examples for which the uncentered account would, against appearances, predict a disagreement:

1: A: “Water is H₂O.” B: “Water isn’t water!”

2: (In a situation in which it is raining), A: “It isn’t raining.” B: “It is raining if and only if it is actually raining!”

in terms of the content of the involved utterances, but the contents of the associated beliefs, makes the right prediction. This account also gives the right assessment of *Disagree*. Here, the associated beliefs of A are: *I have Groat's disease* and *I don't have Groat's disease*. In this case, the associated beliefs are incompatible. (For B, the corresponding incompatible beliefs are: *my addressee has Groat's disease* and *my addressee doesn't have Groat's disease*).

This suggests the following definition of centered disagreement:

Centered Disagreement

Two utterances u and v are in disagreement iff the belief a competent subject with accurate beliefs about u 's and v 's contextual features associates with u is incompatible with the belief she associates with v .

There are several open questions concerning this definition. However, my aim here isn't to give the ultimate characterization of disagreement in centered terms. Rather, I have tried to make the case that the centered account of disagreement is a promising route to explore.

7 The Expressibility of Centered Content

Finally, I want to look at an attempt to reconcile the centered belief account and the FedEx model by rejecting the *Expressibility of Centered Content*. The problem for the FedEx model arises from centered *utterance content*. The Lewisian account of belief only commits us to centered *belief content*. Perhaps a clash between the two can be avoided by keeping utterance content uncentered. But won't the *mind-to-speech principle* fail, if belief content is centered, while utterance content remains uncentered? Not necessarily. Even though some beliefs have centered content, we also have many uncentered beliefs. One may then try to avoid the conflict by restricting the set of beliefs that are expressible: we can only express uncentered beliefs (Egan, 2007), (Moss, 2010). According to that suggestion, even indexical utterances merely express uncentered beliefs. When A utters "I have Groat's disease" she might e.g. express the uncentered belief that *A has Groat's disease*. Since both contents are uncentered, the FedEx model is upheld.

One may doubt whether we can always and in a systematic way find a plausible uncentered candidate for the belief that gets expressed by an indexical utterance. For instance, in the above case A might not believe that she is A, and therefore not have the belief that *A has Groat's disease*. Moreover, there is a rather straightforward argument which seems to lead us from centered belief content to centered utterance content. This argument seems effective against most proposals of assigning uncentered content to utterances, e.g. it seems to rule out the identification of utterance content with the sentence's semantic value on any standard semantic theory. However, I am going to present a way to counter this argument and avoid commitment to centered utterance content, which may also answer the initial worry. Yet, this option is only feasible by taking on further theoretical commitments, which to me seem rather costly—the Recentring model offers a way to avoid these costs.

The argument for centered utterance content focusses on the receiver side of communication. It is a datum that we can learn centered information from utterances of ordinary English sen-

tences (at least, once we have granted that there is centered belief content). For instance, if I don't know what time it is, you can easily resolve my ignorance by uttering "It is noon". Trusting your utterance, I will acquire the centered temporal information that it is noon. How can we explain that I come to learn this centered information? The obvious answer is that the content of the utterance itself is centered.²⁰ This suggests the following argument:

Centered Utterance Content

1. Competent hearers can acquire specific centered information from certain utterances.
2. One can acquire this centered information only if the content of these utterances is itself centered.
3. *Conclusion*: There are utterances whose content is centered.²¹

We have seen that the first premise of the argument is uncontroversial. The problematic premise is the second one. Is it really true that the only way to explain how one can acquire the relevant centered information from an utterance is to assume that the utterance has centered content? Or is there an alternative explanation in terms of uncentered utterance content?

It is difficult to see how one could derive centered information, i.e. information about one's location *within* a world, from uncentered information, i.e. information that distinguishes only *among* worlds.²² Yet, someone who already possesses centered information can learn a further piece of centered information (in part) from uncentered information. In interpreting an utterance, we can typically rely on a wealth of centered background information. Using that information, we are able to derive new centered information from uncentered information. Assume, for instance, that S's belief state is characterized by the following two centered worlds: $\langle w_1, t_1, x_1 \rangle$, $\langle w_2, t_2, x_2 \rangle$. If S now acquires *uncentered* information that excludes w_1 , S learns the *centered* information that she is x_2 in w_2 at t_2 .

That we can acquire centered from uncentered information together with centered background information presents a problem for premise 2. We can now see that there is a potential alternative explanation for how we acquire centered information from utterances, which merely employs *uncentered* utterance content. According to this alternative, the hearer possesses centered background information and then combines this with the uncentered utterance content to arrive at the centered information in question.

There are two challenges for this alternative explanation:

1. How can we guarantee that the hearer reliably possesses the right kind of centered background information?

²⁰Note that we aren't presupposing that the utterance content is the same as the acquired centered content. The explanation in centered terms can then be the one provided by the *Recentering model*.

²¹An argument along these lines can be found in (Torre, 2010, pp. 101–103) and (Ninan, 2010a, §3).

²²It isn't the case that uncentered propositions never imply centered ones. Firstly, the uncentered proposition p entails the centered proposition p or c , where c is a centered proposition. However, the centered information in question is typically not of that form. Secondly, the uncentered proposition that *everything is F* entails the centered proposition that *I am F*, that *you are F*, etc. This will not help either, since whenever the general uncentered information is available, the entailed centered information isn't discriminatory at all, whereas the information we acquire in communication often is.

2. Is it not possible that there are subjects who cannot avail themselves of any centered information, but can nevertheless acquire the right centered information from an utterance?

Firstly, an explanation in terms of uncentered utterance content would not be adequate, if it just applied to cases in which the hearer happens to have the right kind of centered background beliefs. We want a story about how the required centered information is *systematically* available to the hearer, i.e. which covers all cases of successful communication. Secondly, assume that Lewis' two gods don't possess any centered information whatsoever. Still, can we not just tell them where they are? The explanation as to why we cannot do that will show that there may be a way to answer both challenges at one fell swoop.

Linguistic communication crucially involves perceivable vehicles of communication.²³ However, neither god perceives the world from a specific perspective. If one were to try to tell one of the gods where he is, he would merely know that a certain utterance event has taken place at a certain location. Neither of them would be able to figure out whether or not he is the addressee of the utterance. While the two characters of Lewis' story lack the centered information that relates them to the utterance, we aren't so unfortunate. In actual cases of communication, the hearer perceives a given utterance event from a certain point of view. She thereby acquires a specific kind of *centered perceptual content*, something like: *utterance "u" is presently occurring over here*. Hence, her perception generally makes available to the hearer *centered* information which relates her to the utterance token. Relying on this perceptual content, we get a guarantee that every hearer can avail herself of centered information, which is systematically connected to the utterance. A promising alternative explanation in terms of uncentered utterance content will then utilize the centered perceptual content about the utterance token as the relevant centered background information.

Next, we have to find a systematic assignment of uncentered utterance content that together with this information yields the sought after centered information acquired by the hearer. On the present proposal, it seems that the uncentered utterance content will have to involve the utterance token as well.²⁴ Here is an example that illustrates the basic strategy: A doesn't know what time it is. B utters: "It is noon". A acquires the centered information that it is noon. Both perceptual content and utterance content involve the utterance token. We can roughly paraphrase the *centered* perceptual content as: *utterance "u" is occurring now*, while the corresponding *uncentered* utterance content can be put thus: *utterance "u" is occurring at noon*. The hearer can then combine the perceptual content with the utterance content and arrive at the sought after centered information: *it is noon now*.

It is easy to show that, in principle, there is an uncentered proposition which together with the right centered proposition contains the acquired centered information that we are after. Take, for instance, the *uncentered singular* proposition about the utterance token that it occurs at noon. Combine this with the *centered singular* proposition that the utterance token is occurring now, and you'll obtain the right centered information. Alternatively, we can employ corresponding

²³It is conceivable that there are other kinds of communication, which don't involve such vehicles, e.g. telepathy. The following doesn't apply to these.

²⁴Of course, I cannot completely exclude the possibility that I may have overlooked an alternative assignment of centered background information and uncentered utterance content, which could also cover all relevant cases.

qualitative propositions which refer to the utterance token by some individuating property.

The existence of these propositions shows that it is in principle possible to derive the relevant centered information in part from uncentered information. However, a plausible explanation of what happens in actual cases of communication then has to include the claim that the relevant propositions, be they singular or general, are in fact what we express by our utterances and what we acquire in perceiving an utterance. This is where things become problematic.

Let us first consider the option of using qualitative propositions. One possibility for identifying the utterance token is by its spatio-temporal position. This option isn't promising: we cannot ascertain the spatio-temporal position of an utterance simply by perceiving it. Moreover, the relevant spatio-temporal information is often exactly the information the hearer is lacking and the speaker is trying to provide.

Alternatively, we can pick out the utterance token by its perceivable qualitative features. According to that proposal, the uncentered utterance content is something like: *there is an utterance with perceivable features U and that utterance is occurring at noon*, while the complementary perceptual centered proposition can be put like this: *there is an utterance with perceivable features U and that utterance is occurring now*. This option doesn't work either, because the ascribed utterance content is too weak. Assume that there is a qualitatively identical utterance, i.e. one that shares all the perceivable features U and that this utterance is made at 11 o'clock. According to the above suggestion, the truth condition of this utterance are fulfilled: even at 11 o'clock it is true that there is an utterance with perceivable features U which occurs at noon, namely the original utterance. Obviously, we cannot truthfully utter "It is noon" at 11 o'clock.

The attempted explanation in terms of qualitative propositions is unsuccessful. What about employing singular propositions? We can give the utterance token "It is noon" a name and call it "Hans". The content of the utterance can then be understood as the *uncentered* singular proposition: *Hans is occurring at noon*, and the coordinated perceptual content as the *centered* singular proposition: *Hans is occurring now*. From these two singular propositions, the hearer can infer the centered information: *it is noon now*.

This proposal, too, faces difficulties. Within modal accounts of content, it is problematic to characterize doxastic states by singular propositions (Lewis, 1981). Believing a singular proposition about an object is to have very strong beliefs about that object. If I accept the singular proposition that *a is F*, then in all my doxastic alternatives *a* will exist and have property *F*. However, in all my doxastic alternatives *a* will also possess its essential properties. Hence, ascribing a singular proposition about *a* to a subject is to ascribe to that subject true beliefs about *a*'s essential properties. Yet, there aren't many cases where we actually possess this information.²⁵ This problem equally applies to the above proposal: plausibly, I can come to learn the information conveyed by the utterance, without having any beliefs about the essential properties of the utterance token used. The same goes for my perception of the utterance.

There are ways to avoid this problem. One might adopt an externalist perspective according to which the propositions ascribed to a subject aren't intended to capture that subject's perspective on the world, but rather describe her belief system from the theorist's external point of view

²⁵To be clear, I am not questioning the existence of singular (i.e. *de re*) beliefs about objects, but merely that they are best understood as beliefs in singular propositions.

(Stalnaker, 1999, p. 25). One might separate (merely) epistemic from metaphysical possibilities (Soames, 2006). Lastly, one might deny that objects have non-trivial essences. To me, none of these options seem very attractive.

There is a further worry. Not only may the hearer lack beliefs about the utterance's essential properties, she may likewise lack individuating information about which utterance token she just encountered. The hearer can know that the utterance is true, without being able to exclude situations in which a qualitatively identical, but numerically different utterance token, call it *Franz*, was being produced. However, the above proposal seems to demand that the hearer knows the identity of the token. Adopting an externalist perspective may provide a way out of this problem as well. From an externalist point of view, the hearer's causal rapport with the utterance token may be sufficient to license an ascription of the above singular propositions.

Even though explaining the acquisition of centered information with the help of uncentered utterance content is a theoretical option, it seems problematic. Yet, since sticking to uncentered utterance content has the advantage of reconciling the FedEx model with the centered belief account, it may, all things considered, be a bullet worth biting.

The Recentering model provides an explanation which avoids the contentious commitments. Why is this model better off? For one thing, it doesn't demand that the utterance content involves the utterance token. Furthermore, even though the model does require the hearer to make reference to the utterance, it endows her with additional resources for picking it out. She can take advantage of self-locating descriptions, such as *the utterance token I am presently perceiving*.²⁶ We can then reformulate the initial argument for centered utterance content as an inference to the best explanation: the assumption that there is centered utterance content provides the most plausible explanation of how we acquire centered information in communication. And the Recentering model is a promising way to accommodate the centered content of utterances.

8 Conclusion

Many of our beliefs represent the world from a specific location within it. Our language seems to allow us to express this perspectival aspect of thought—we aren't restricted to only giving voice to beliefs that represent the world from a god's eye view. The Recentering model of communication explains how this is possible. As the speaker doesn't merely characterize the world she shares with the hearer, but rather represents something about her individual location within it, the hearer cannot directly endorse the information expressed. She gains information about her own location from information about a different location and her beliefs about how she is related to that location. On this picture, it isn't the case that there is always a single piece of information which travels from speaker to hearer. We have to give up the FedEx model. Insofar as the Recentering model is a viable alternative, this may be a price worth paying.²⁷

²⁶For cases in which the hearer encounters several utterance tokens at the same time, she can pick out a single token by its position within her sensory field, or as the one that is the focus of her attention.

²⁷For helpful comments and discussion, I would like to thank Jochen Briesen, David Chalmers, John Cusbert, Andy Egan, Edward Elliott, Peter Fritz, Frank Jackson, Dirk Kindermann, Jens Kipper, Dilip Ninan, Daniel Nolan, Peter Pagin, Josh Parsons, Jonathan Schaffer, Wolfgang Schwarz, an anonymous referee, as well as audiences at

References

- Chalmers, D.: The Component of Content. In: D. Chalmers (ed.) *Philosophy of Mind: Classical and Contemporary Readings*, pp. 608-633. Oxford University Press (2002)
- Chalmers, D.: Epistemic two-dimensional semantics. *Philosophical Studies* **118**, 153–226 (2004)
- Egan, A.: Epistemic modals, relativism and assertion. *Philosophical Studies* **133**, 1–22 (2007)
- Feit, N.: *Belief about the Self: A Defense of the Property Theory of Content*. Oxford University Press (2008)
- Heim, I.: Lecture notes on indexicality (2004). MS
- Jackson, F.: *From Metaphysics to Ethics: A Defence of Conceptual Analysis*. Oxford University Press (1998).
- Jackson, F.: *Language, Names, and Information*. Blackwell, Oxford (2010)
- Kaplan, D.: Demonstratives. In: J. Almog, J. Perry, H. Wettstein (eds.) *Themes from Kaplan*. Oxford University Press, New York (1989)
- King, J.C.: *The Nature and Structure of Content*. Oxford University Press (2007)
- Lewis, D.: Attitudes *de dicto* and *de se*. *The Philosophical Review* **88**, 513–43 (1979)
- Lewis, D.: Index, context, and content. In: S. Kanger, S. Öhman (eds.) *Philosophy and Grammar*, pp. 79–100. D. Reidel, Dordrecht (1980)
- Lewis, D.: What puzzling Pierre does not believe. *Australasian Journal of Philosophy* **59**, 283–89 (1981)
- Lewis, D.: Reduction of mind. In: S. Guttenplan (ed.) *The Blackwell Companion to the Philosophy of Mind*, pp. 412–31. Oxford University Press, Oxford (1995)
- Meacham, C.: Unravelling the tangled web: Continuity, internalism, non-uniqueness and self-locating beliefs. In: T. Szabó-Gendler, J. Hawthorne (eds.) *Oxford Studies in Epistemology*, vol. 3. Oxford University Press, Oxford (2010)
- Moss, S.: Updating as communication. *Philosophy and Phenomenological Research*, DOI: 10.1111/j.1933-1592.2011.00572.x.
- Ninan, D.: *Imagination, content, and the self*. Ph.D. thesis, MIT (2008). URL <http://www.st-andrews.ac.uk/~djn1/diss.pdf>
- Ninan, D.: De se attitudes: Ascription and communication. *Philosophical Compass* **5**(7), 551–567 (2010)

seminars in Canberra, Cologne, Dunedin, and Stockholm.

- Ninan, D.: Self-location and other-location (2010). URL <http://www.st-andrews.ac.uk/djn1/otherlocation.pdf>. MS
- Nolan, D.: Selfless desires. *Philosophical and Phenomenological Research* **73**(3), 665–679 (2006)
- Perry, J.: The problem of the essential indexical. *Nous* **13**, 3–21 (1979)
- Schwarz, W.: Changing minds in a changing world. *Philosophical Studies*, DOI:10.1007/s1109801196990.
- Soames, S.: Philosophical implications of the Kripkean necessary aposteriori. *Philosophical Issues* **16**, 288–309 (2006)
- Stalnaker, R.C.: Assertion. In: P. Cole (ed.) *Syntax and Semantics Vol. 9: Pragmatics*. Academic Press, New York (1978)
- Stalnaker, R.C.: Indexical belief. *Synthese* **49**, 129–151 (1981)
- Stalnaker, R.C.: Belief attribution and context. In: R. Grimm, D. Merrill (eds.) *Contents of Thought*, pp. 140–56. University of Arizona Press, Tucson (1988)
- Stalnaker, R.C.: *Context and Content: Essays on Intensionality Speech and Thought*. Oxford University Press, Oxford (1999)
- Stalnaker, R.C.: *Our Knowledge of the Internal World*. Oxford University Press, New York (2008)
- Torre, S.: Centered assertion. *Philosophical Studies* **150**, 97–114 (2010)