

Self-Location and Other-Location*

Dilip Ninan
dilip.ninan@tufts.edu
May 6, 2013

1 Introduction

According to one tradition in the philosophy of language and mind, the content of a psychological attitude can be characterized by a set of possibilities. On the classic version of this account, advocated by Hintikka (1962) and Stalnaker (1984) among others, the possibilities in question are possible worlds, ways the universe might be.¹ Lewis (1979, 1983a) proposed an alternative to this account, according to which the possibilities in question are possible individuals or *centered worlds*, ways an individual might be.² The motivation for the centered worlds theory has primarily to do with self-locating – or *de se* – attitudes. The focus of this paper is on the less-discussed question of how other-locating – or *de re* – attitudes ought to be treated within this framework. Most advocates of what we might call the *modal approach* to attitudes, Stalnaker and Lewis included, offer some kind of descriptivist solution to the well-known problems that other-locating attitudes raise. There are intramural differences between Stalnaker, Lewis, and other modal theorists (e.g. *two-dimensionalists*) on a number of issues: on the precise nature of the descriptivism involved, how attitude content relates to the asserted content of the sentences we utter, and on the proper semantic treatment of attitude reports. I pass over these differences to focus on a problem common to these various approaches: all face a problem when it comes to characterizing the contents of *counterfactual attitudes* like imagining, dreaming, and wishing.

I then show how the problem can be solved if we adopt a different account of attitudinal alternatives, the possibilities used to characterize the content of an attitude. Rather than using centered worlds, I propose to use *multi-centered worlds*, which one can think of as a way a group of individuals might be. I discuss two versions of this proposal. On the first, *de se* contents are what

*For helpful comments, thanks to Derek Ball, Josh Dever, Andy Egan, Gabriel Greenberg, Richard Holton, Dirk Kindermann, Sarah Moss, François Recanati, Robert Stalnaker, Will Starr, Seth Yalcin, and two anonymous referees for *Philosophy and Phenomenological Research*.

¹See also the papers in Stalnaker (1999). In more recent work, Stalnaker has come to accept a version of the centered worlds approach (Stalnaker 2008, Ch. 3). I thus concentrate on that approach, since the pure possible worlds theory now appears to have few advocates.

²See also Lewis (1986) and Lewis (1995).

Perry (1979) calls “propositions of limited accessibility,” i.e. propositions that only a single agent can entertain. On the second, *de se* contents are *relativistic* as they are on the centered worlds approach. I close with a tentative discussion of how we might try to decide between these two versions of the proposal.

2 Self-location

Although my primary interest concerns other-locating attitudes, I begin with self-locating attitudes, since it is these that have been taken to motivate the centered worlds theory. There are a number of different arguments that Lewis and others have offered in favor of centered worlds, but I will focus on one that concerns the role of *de se* attitudes in the explanation of action.³

Building on earlier work by John Perry,⁴ Lewis observes that two agents who share an appropriate set of *de se* beliefs and desires will, *ceteris paribus*, be disposed to behave in the same way. Suppose, for example, that you and I are both in a belief state that we could each express by saying, “I am being attacked by a bear” and that both of us believe that the best way to avert a bear attack is to curl up into a ball and remain still. Then we will, other things being equal, both be disposed to curl up and remain still. On Lewis’s view, the point of assigning content to attitudes is to characterize their causal role vis-à-vis stimuli, behavior, and other attitudes. Since your belief that you will be attacked by a bear and my belief that I will be so attacked have the same causal role in our respective cognitive systems, they ought to be assigned the same content.

There is much more to say about this last claim, and I will return to it later in the paper (see §8).⁵ But if we *do* we accept Lewis’s premise for the moment, we will be forced to admit that contents have *relative truth values*: one and the same content can be true for one person, false for another. To see this, let p be the content that you and I both believe. Now suppose that I am in fact being attacked by a bear, but you are not (it’s merely a man in a bear suit). Then my belief is true and yours is false. But a belief is true (false) only if the content of the belief is true (false). Since my belief is true, and the content of my belief is p , p must be true. Since your belief is false, and the content of your belief is p , p must be false. But p can’t be both true and false, so something has gone wrong.

Lewis’s solution is to give up the idea that p is true or false *simpliciter*; rather it is something that can be true relative to one agent, false relative to

³One can find something similar to this argument in Lewis’s discussion of the Hume-Heimson case (Lewis 1979, 142-143). See also the introduction to Stalnaker (1999), Egan (2010), and Moss (2012). For other arguments for the centered worlds theory, see Lewis (1979), Lewis (1983a), and Ninan (2012).

⁴See Perry (1977, 1979).

⁵Perry (1979) denies that you and I believe the same content. He captures the psychological similarity between us by saying that we are in the same *belief state*, and he thinks belief states in this sense can be characterized by centered (“relativized”) propositions. It is not entirely clear whether this difference between Lewis and Perry is substantive or verbal (Lewis 1979, 151-152). See Perry (2001) for Perry’s more recent view.

another. Since my belief is true, p is true for me; since yours is false, p is false for you. Problem avoided. But what kind of thing is p ? Lewis accepts the modal approach to attitudes, so p should be a set of possibilities. But it cannot be a set of possible worlds, for the following reason. A set of possible worlds (a possible worlds proposition) q is true at a world w' just in case w' is an element of q . You and I inhabit the same world, w . So if p is a possible worlds proposition, p is either true at w or false at w . If it is true at w , then it cannot be the content of your belief, since your belief is false; if false at w , it cannot be the content of my belief, since my belief is true.

Lewis's alternative is to take p to be a set of *centered worlds*, where a centered world is a triple consisting of a world, a time, and an individual who exists at the time and world in question.⁶ This proposal requires some explanation.

The key notion of the centered worlds proposal is what it is for a centered world to be compatible with what an agent believes. In his papers on this topic, Lewis essentially treats this as a primitive notion, not one that gets characterized in independent terms. But although it is a primitive notion, we can say various things to convey its intuitive content. For example, we might say the following, for any agent x , time t , and world w :

- (1) A centered world (w', t', x') is compatible with what x believes at t in w iff x' has, at t' in w' , all of the properties that x believes *de se* (at t in w) that she herself has.

Notions of compatibility for other attitudes (desire, imagination, etc.) can be given parallel characterizations.

A centered proposition is a set of centered worlds, and an agent believes a centered proposition q just in case all of her centered belief worlds – all of her doxastic alternatives – are contained in q . If an agent x believes centered proposition q at time t in world w , her belief is true just in case (w, t, x) is in q .

Note that our gloss on the notion of compatibility employs the notion of belief *de se*. This underscores the fact that what we have here is not a reductive theory of *de se* attitudes, but rather a proposal for how to characterize their content. Reduction may come at a later stage, when we ask in virtue of what an agent should count as believing or desiring a centered proposition.

Now let us apply all this to the bear attack case. I believe, at t in w , that I am being chased by a bear. So every centered world (w', t', x') compatible with what I believe at t in w is such that x' is being chased by a bear at t' in w' . So I believe p :

$$p = \{(w', t', x') : x' \text{ is being chased by a bear at } t' \text{ in } w'\}$$

Since I am being chased by a bear, (w, t, me) is in p , and so my belief is true. You also believe, at t in w , that you are being chased by a bear. So every

⁶Actually, since Lewis takes individuals to be world-bound and believes in temporal parts, he identifies p with a set of world-bound time-slices, individuals that determine a world and a time. Using centered worlds allows us to consider the same proposal within alternative metaphysical frameworks.

centered world (w', t', x') compatible with what you believe at t in w is such that x' is being chased by a bear at t' in w' . So you believe p as well. But since you are not being chased by a bear, (w, t, you) is not in p , and so your belief is false. Thus, we believe the same thing, and that thing is true for me, but false for you.

3 Other-location

How do we treat other-locating – or *de re* – attitudes on the centered worlds approach? We can bring out a potential problem for the approach by considering a well-known case due to Quine (1956):

There is a certain man in a brown hat whom Ralph has seen under questionable circumstances; suffice it to say that Ralph suspects that he is a spy. There is also a gray-haired man who Ralph sees at the beach one day; Ralph recognizes this man to be Bernard J. Ortcutt, the town mayor, and Ralph believes that this man is no spy. Now Ralph does not know it, but the men are one and the same.

What does the centered worlds theorist say about this case and others like it: Pierre who fails to realize that London is Londres, the astronomer who has yet to discover that Hesperus is Phosphorus, etc.?

Let t and w be the time and world (respectively) of Ralph's belief. Since Ralph believes (at t in w) that the man in the brown hat is a spy, we might be tempted to say that each of the centered worlds (w', t', x') compatible with what Ralph believes (at t in w) is such that this man is a spy at t' in w' . And since Ralph believes (at t in w) that Bernard J. Ortcutt is no spy, we might want to say that each of the centered worlds (w', t', x') compatible with what Ralph believes (at t in w) is such that Ortcutt is not a spy at t' in w' . But since the man in the brown hat just is Ortcutt, if we say both these things, then we must admit that no centered worlds are compatible with what Ralph believes at t in w , since no centered worlds (w', t', x') are such that that man/Ortcutt both is and is not a spy at t' in w' .

The problem with saying that no centered worlds are compatible with what Ralph believes in this scenario is that Ralph seems to have a coherent picture of what the world is like. Indeed, it seems that we have no trouble conceiving of what things are like from Ralph's perspective: there is a man in a brown hat who Ralph saw under questionable circumstances, and this man is a spy; there is another man, different from the first, who has gray hair and who Ralph saw on the beach, and this man is not a spy. The present approach to other-location doesn't seem to capture Ralph's intuitive picture of the world.

Note that this problem is not a problem for the centered worlds proposal *per se*: exactly the same problem would arise if we were using (uncentered) possible worlds to represent Ralph's doxastic alternatives instead. The general problem here is a problem for the idea that the content of a *de re* belief to the effect that y is F is the set of possible/centered worlds in which y is F . It seems that, on

a modal theory of attitudes, the content of a *de re* attitude cannot in general be a singular proposition.

Most advocates of the modal approach to attitudes accept some kind of descriptivist resolution to this problem. For example, the approach taken by Stalnaker is of essentially this kind – this is quite clear in Stalnaker (2008, Ch. 3).⁷ The same goes for two-dimensionalists like David Chalmers and Frank Jackson.⁸ But rather than trying to describe this descriptivist approach in very general terms, I will, for the sake of concreteness, focus on Lewis’s particular version of the descriptivist solution. But the problem I will later raise for Lewis’s theory should be easy to extend to the other versions of this general approach.

Lewis’s account of other-location is a variation on the account of the *de re* pioneered by Quine (1956) and Kaplan (1968). Lewis starts with the notion of a relation of acquaintance, a relation “of a sort apt for the reliable transmission of information” Lewis (1979, 155). Note that Ralph bears two different relations of acquaintance to Ortcutt: on one occasion, he sees Ortcutt wearing a brown hat; on another occasion, he sees gray-haired Ortcutt on a beach. Let Q be the relation that x bears to y just in case y is the unique individual that x has seen wearing a brown hat. Let S be the relation that x bears to y just in case y is the unique gray-haired individual that x has seen on the beach. So Ralph bears both Q and S to Ortcutt.⁹

On Lewis’s account, the basic notion of other-locating or *de re* belief is the notion of believing something about someone *relative to an acquaintance relation*.¹⁰ I take this notion to have reasonably clear intuitive content. Ralph

⁷Stalnaker’s method of *diagonalization* is not an alternative to the descriptivist approach, but a way of repairing a mismatch between the descriptive beliefs intuitively expressed or reported by certain utterances and the normal semantic content of the sentences uttered.

Note that Stalnaker isn’t a descriptivist about other-locating attitudes *per se*. On Stalnaker’s view, when identity confusion is not an issue, an agent can be said to believe a singular proposition. But Stalnaker does adopt a descriptivist solution to cases like the Ortcutt case.

⁸See, for example, Jackson (1998) and Chalmers (2002). (Chalmers might resist being described as a descriptivist, but this subtlety is irrelevant here.)

⁹It may be unrealistic to suppose, e.g., that Ralph has only ever seen one man in a brown hat before. But this simplification is, I think, harmless in the present context.

It is also worth noting that although Lewis takes acquaintance relations to be causal relations of some sort, his solution to the Ortcutt case does not actually require us to understand the term “acquaintance relation” in any particularly robust way. For example, we could, if we wanted, allow the (degenerate) relation that x bears to y just in case y is the shortest spy to count as an acquaintance relation. This may be an abuse of terminology, in which case we can put the point by saying that the relations in question need not be acquaintance relations at all. Whether or not *some* restriction ought to be placed on the class of relevant relations is an issue we can leave open; it seems to me that the issues discussed in this paper are largely separable from the question of how an agent must be related to an individual in order to have a *de re* thought about that individual.

¹⁰Lewis (1979, 152ff.) talks about “ascribing a property to an individual under a description.” But he takes a ‘description’ to be a relation, and in the case of *de re* belief, he requires the relation to be a relation of acquaintance. So this comes to the same thing as my ‘believing something about someone relative to an acquaintance relation.’ But this is, in any case, something of a simplification; something more general is needed for the multiply *de re* case, as when Ralph believes that he is standing in between Jones and Ortcutt. The general notion is something like ‘believing something about a plurality of individuals-relative-to-acquaintance-relations.’

is acquainted with Ortcutt in two different ways, but fails to realize this. He bears relation Q to Ortcutt, and relative to this relation, Ralph believes that Ortcutt is a spy. (We might suppose that, upon seeing the man in the brown hat, Ralph points and whispers, “That man is a spy.”) Ralph also bears relation S to Ortcutt, and relative to this relation, Ralph believes that Ortcutt is not a spy. (We might imagine Ralph nodding as we point to the gray-haired man on the beach and say, “That Bernard Ortcutt is no spy.”) It will often be useful to put the “relative to” qualification next to the referring expression at issue; for example, I will often say things like: “Ralph believes that Ortcutt (relative to Q) is a spy, and believes that Ortcutt (relative to S) is not a spy.” This is somewhat inelegant, but serves the purposes of clarity.

Lewis then offers the following analysis of the intuitive notion of believing something about someone relative to an acquaintance relation:

- (2) An agent x believes, at t in w , that y is F , relative to acquaintance relation R iff:
- (i) x bears R uniquely to y at t in w , and
 - (ii) x believes *de se* (at t in w) that the thing to which he bears R is F .¹¹

Given Lewis’s account of *de se* belief, (ii) says that all the of centered worlds (w', t', x') compatible with what x believes at t in w are such that the thing to which x' bears R at t' in w' is F at t' in w' .

Note that, on this approach, there will be centered worlds compatible with Ralph believes in Quine’s case. At time t in world w , Ralph believes that Ortcutt (relative to Q) is a spy. On Lewis’s account, this means that Ralph believes (at t in w) that the man he saw in the brown hat is a spy. So every centered world (w', t', x') compatible with what Ralph believes (at t in w) is such that the man x' has seen in a brown hat at t' in w' is a spy at t' in w' .¹² Ralph also believes (at t in w) that Ortcutt (relative to S) is not a spy. So he believes (at t in w) that the gray-haired man he saw on the beach is not a spy. So every centered world (w', t', x') compatible with what Ralph believes (at t in w) is such that the gray-haired man that x' saw on the beach at t' in w' is not a spy. But there is no conflict between these two claims about Ralph’s doxastic alternatives. Each of the centered worlds (w', t', x') compatible with what Ralph believes at t in w contains two relevant individuals, both of whom correspond, in some sense, to Ortcutt. One of them will be the person whom x' saw in a brown hat under questionable circumstances in w' , and this individual will be a spy in w' ; the other will be a gray-haired person that x' sees on the beach in w' , and this person will not be a spy in w' . Since the men in question are distinct, there is no contradiction in saying that the one is a spy and the other is not.

¹¹This is Lewis’s account of the singly *de re*. One would need something more general to handle the multiply *de re*.

¹²I am being a bit sloppy about the temporal aspects of Ralph’s beliefs, but this will not affect any of the issues discussed in the paper.

Given this account of *de re* belief, an agent x believes *de se* at t in w that she is F just in case she believes *de re* that she is F relative to the relation I , identity (Lewis 1979, 156-157). This is because x believes, at t in w , that x is F relative to I iff:

- (i) x bears I to x at t in w , and
- (ii) all the centered worlds (w', t', x') compatible with what she believes (at t in w) are such that the thing to which x' bears I at t' in w' is F at t' in w' .

But since I is identity, (i) is satisfied so long as x exists at t in w , and so is entailed by (ii). And (ii) is equivalent to the claim that all the centered worlds (w', t', x') compatible with what x believes (at t in w) are such that x' is F at t' in w' . Since the latter is just what it means, on the centered worlds account, to say that x believes *de se* at t in w that she is F , belief *de se* is just belief *de re* where the *res* is the agent and the relation of acquaintance is I , identity.

4 Counterfactual attitudes

Up until now, our discussion has mostly focussed on the attitude of belief. But of course we want an account that applies to all the other content-bearing mental states: hopes, fears, desires, episodes of imagining, dreams, and so on. But a problem arises when we try to extend Lewis's account of other-location to counterfactual attitudes like imagining, wishing, and dreaming.

Consider imagining. As with belief, we have an intuitive notion of imagining something about someone relative to an acquaintance relation. An example will help to illustrate this. Suppose Ralph sees the gray-haired man on the beach who he recognizes as mayor Ortcutt. Ralph might then have an imagining he could report by saying, "I'm imagining that that man [he points at the man on the beach] is flying a kite." Suppose that Ralph is, at the same time, holding a photograph of the man in the brown hat who he earlier saw behaving suspiciously. (As before, Ralph doesn't realize that the man on the beach is the man in the brown hat.) Note that Ralph might have a different imagining, one he could report by saying, "I'm imagining that this man [he points at the man in the photograph] is flying a kite." Intuitively, these are two different imaginings; imagining is fine-grained in the familiar way that virtually all psychological attitudes are.

In the first case, Ralph imagines something about Ortcutt partly in virtue of his bearing relation S to Ortcutt (the 'see on the beach' relation). So we can say that in the first case, he is imagining that Ortcutt (relative to S) is flying a kite. In the second case, Ralph imagines something about Ortcutt partly in virtue of his bearing relation Q to Ortcutt (the 'see in a brown hat' relation). So we can say that in the second case, he is imagining that Ortcutt (relative to Q) is flying a kite. Here we are simply extending our way of talking about *de re* belief to *de re* imagining.

The natural extension of Lewis's analysis of belief *de re* to imagination *de re* is this:

- (3) An agent x imagines, at t in w , that y is F , relative to acquaintance relation R iff:
- (i) x bears R uniquely to y at t in w , and
 - (ii) x imagines *de se* (at t in w) that the thing to which he bears R is F .¹³

Again, (ii) breaks down to this: every centered world (w', t', x') compatible with what x imagines at t in w is such the individual to whom x' bears R at t' in w' is F at t' in w' . Note that this means that the definite description in (ii) is understood to have a narrow-scope reading.

But there are two problems with (3). Claim (3) should be understood as a universally quantified claim, instances of which are biconditionals. And we can find instances of it that fail in the left-to-right direction, and instances of it that fail in the right-to-left direction. To see an instance of the first kind, let's suppose that Ralph is watching the brown-hatted Ortcutt behaving suspiciously. Notice that, of the various things Ralph might imagine about Ortcutt (relative to Q , the 'see in a brown hat' relation), some of them are compatible with Ralph's beliefs about Ortcutt (relative to Q), while some are not. For an example of the latter, note that although Ralph believes that Ortcutt (relative to Q) is a spy, Ralph might imagine a scenario in which Ortcutt (relative to Q) never took up espionage. That is, Ralph might point to the brown-hatted man he is watching and say, "I am imagining a scenario in which he is not a spy." The centered worlds account has no trouble characterizing the content of this imagining; it is simply the set of centered worlds in which the individual to whom the center bears Q – the individual who the centers sees wearing a brown hat – is not a spy.

But the centered worlds account runs into a problem in cases in which what Ralph imagines about Ortcutt (relative to Q) conflicts with his bearing relation Q to Ortcutt. For example, suppose Ralph is again looking at the brown-hatted man before him, and has an imagining he could report by saying:

"I'm imagining a situation in which I never laid eyes on that man [Ralph points at the brown-hatted man before him]. In my imaginary scenario, that man is a recluse who few people have ever seen."

¹³I emphasize that we are considering an *extension* of Lewis's account. Lewis himself says this at one point:

If my theses hold for belief and knowledge and desire, then also they hold for any attitude which amounts to lack of belief or knowledge or desire... I'm not sure anything is left out—perhaps some ill-understood attitudes of imagining, conceiving, contemplating, or entertaining a thought. (Lewis 1979, 145)

But (as we shall see) our problem arises not just for imagining, conceiving etc., but also for the desire-like attitude of wishing. (Thanks to Richard Holton and Will Starr for reminding of this passage from Lewis's paper.)

Now it is intuitively true that Ralph is imagining that Ortcutt (relative to Q) is a recluse whom he has never seen. But according to (3), this is true only if two things hold (assume t and w are the time and world of the imagining):

- (i) Ralph bears Q uniquely to Ortcutt at t in w , and
- (ii) all the centered worlds (w', t', x') compatible with what Ralph imagines at t in w are such that there is a unique individual y' to whom x' bears Q at t' in w' , and y' is a recluse whom x' has never seen at t' in w' .

Claim (i) is true, since Ortcutt is the man Ralph has seen in a brown hat at t in w , and Q is the ‘see in a brown hat’ relation. But (ii) is false. To see this, first note that since Ralph’s imagining is intuitively coherent, there should be at least one centered world compatible with what he imagines; this means that (ii) is not vacuously true. Now recall that Q is the relation that x bears to y just in case y is the unique individual that x sees wearing a brown hat. And note that there is no centered world (w', t', x') in which there is an individual y' such that: (a) y' is the unique individual who x' sees wearing a brown hat at t' in w' , and (b) y' is a recluse whom x' has never seen at t' in w' . There is no such centered world because y' cannot both have been seen in a brown hat by x' and never have been seen by x' . Since there are no centered worlds that meet this condition, it cannot be that all the centered worlds compatible with what Ralph imagines at t in w meet this condition. Hence (ii) is false, and the relevant biconditional (the relevant instance of (3)) fails in the left-to-right direction.¹⁴

I’ll return to consider some responses to this objection in a moment, but first I want to explain why I think some instances of (3) fail in the right-to-left direction. Suppose Ralph is again watching Ortcutt behaving suspiciously in his brown hat. But now Ralph’s thoughts wander away from Ortcutt, and he begins to imagine a scenario in which he sees exactly one man in a brown hat, and in which this man is flying a kite. As I am thinking of this situation, Ralph’s imagining is not about Ortcutt nor any other individual of his acquaintance (other than himself): he is simply imagining that he is watching someone or other wearing a brown hat and flying a kite. Other than the fact that his imaginary scenario contains himself, his imagining is purely general or *de dicto*. This seems clearly possible, and it seems that if t and w are the time and world of his imagining, then it should not be true that Ralph is imagining at t in w that Ortcutt (relative to Q) is flying a kite. This should not be true, since, by stipulation, Ralph is not imagining anything about Ortcutt. But consider this biconditional, which I take to be an instance of (3):

Ralph imagines, at t in w , that he is watching Ortcutt (relative to Q) flying a kite iff:

- (i) Ralph bears Q uniquely to Ortcutt at t in w , and

¹⁴There is obviously a similarity between this objection to the centered worlds theory and Kripke’s modal objection to the description theory of proper names (Kripke 1980).

- (ii) all the centered worlds (w', t', x') compatible with what Ralph imagines at t in w are such that the individual to whom x' bears Q at t' in w' is flying a kite at t' in w' .

As I said, it seems to me that the left-hand side of this biconditional is false. But the right-hand side is true: Ralph bears Q to Ortcutt uniquely, and each of Ralph's centered imagination worlds is such that the individual to whom the center bears Q is flying a kite. Each of Ralph's centered imagination worlds contains exactly one individual who the center sees wearing a brown hat, and that individual is flying a kite. So (3) predicts that Ralph should count as imagining that he is watching Ortcutt (relative to Q) flying a kite. But Ralph is plainly not doing this: his imagining is purely general and is not about the brown-hatted man in front of him.¹⁵

Two points about the scope of these objections. First, it is not just the mental state of imagining that is at issue; other counterfactual attitudes like wishing and dreaming give rise to a similar problem. For example: Ralph might wish that he had never seen Ortcutt (relative to Q). Second, although I have been focussing on Lewis's approach to other-location, the objection would seem to apply quite generally to any modal theory of attitudes that employs some sort of descriptive treatment of the *de re*. (I will have a little bit more to say about this below, when I discuss the possibility of a two-dimensionalist solution to this problem.)

One might think that both of these problems arise only because we are taking a too-simplistic view of the the acquaintance relation relative to which Ralph is thinking of Ortcutt. After all, it is not really plausible that Ralph has only ever seen one man wearing a brown hat. So the relevant acquaintance relation is not really Q , but something richer, something based on Ralph's perceptual experience of Ortcutt. Perhaps it is relation T , the relation x bears to y just in case y is the unique individual that looks ϕ , where ϕ is the way Ortcutt looks to Ralph on the occasion in question.

But this move won't solve the general problem. Take the left-to-right problem again. Ralph could imagine a scenario in which Ortcutt (relative to T) does not look ϕ , perhaps because of an unfortunate accident that he had as a child (he says, "I'm imagining that that man [he points at the brown-hatted man

¹⁵Some people have wondered whether *counterpart theory* (cf. Lewis 1968) could be brought in to resolve these problems. The answer to this question depends in part on what one means by "counterpart theory." Lewis's account of *de re* attitudes can be regarded as a type of counterpart theory (Lewis 1983a); obviously *that* version of counterpart theory is of no help since it is the target of our objections. But the account I shall be proposing in the next section could also be considered a version of counterpart theory, and I believe that that account does solve the problem. Perhaps the important point is this: it is difficult to see how any account on which counterpart relations are *qualitative similarity relations* could work, given the problem just mentioned. For consider two imaginings Ralph might have: he might imagine that there is some man or other possessing certain qualitative characteristics; he might instead imagine that Ortcutt (relative to Q) possesses those very same characteristics. It is hard to see how a qualitative-counterpart-theoretic account could distinguish these two cases. (A modal theory of attitudes that employs non-qualitative counterpart relations faces a different sort of problem; see Ninan (2012) for discussion.)

before him] looks very different from the way he in fact looks”). Or Ralph could imagine a scenario in which Ortcutt (relative to T) dies as a child; presumably the doomed child of Ralph’s imagining looks little like the strapping man before him, and so again is unlikely to look ϕ .

In fact, I think there is a very general reason why this sort of strategy for avoiding the problem is not going to succeed. Whatever the ‘right’ acquaintance relation relative to which Ralph is thinking about Ortcutt is, the fact that Ralph stands in that relation to Ortcutt is likely to be a contingent fact. After all, what typically puts us in a position to think about an object is our standing in some contingent relation to that object. This makes it plausible that, no matter exactly what that relation is in the Ralph-Ortcutt case, Ralph will be able to imagine a scenario in which Ortcutt exists but in which he (Ralph) does not bear that relation to Ortcutt. So the problem cannot be avoided simply by making sure that we’ve selected the ‘right’ acquaintance relation.

Let us put the account of *de re* imagining we are discussing before ourselves again:

- (3) An agent x imagines, at t in w , that y is F , relative to acquaintance relation R iff:
 - (i) x bears R uniquely to y at t in w , and
 - (ii) x imagines *de se* (at t in w) that the thing to which he bears R is F .

I said that the definite description – “the thing to which he [i.e. x] bears R is F ” – is read with narrow scope. But what if we read the same description with wide scope or (what amounts to the same thing) with a tacit “actually” between “he” and “saw”? I take it the proposal here is to replace (3) with (4), or to read (3) as if it expressed (4):

- (4) An agent x imagines, at t in w , that y is F , relative to acquaintance relation R iff:
 - (i) x bears R uniquely to y at t in w , and
 - (ii) every centered world (w', t', x') compatible with what x imagines (at t in w) is such that the thing to which x bears R at t in w is F at t' in w' .

Consider our initial problem case in light of (4). According to this view, Ralph imagines, at t in w , that he never saw Ortcutt (relative to Q) only if every centered world (w', t', x') compatible with what Ralph imagines at t in w is such that x' never sees, at t' in w' , the individual to whom Ralph bears Q at t in w . Now the individual to whom Ralph bears Q at t in w is just Ortcutt. So on this proposal, Ralph imagines, at t in w , that he never saw Ortcutt (relative to Q) only if every centered world (w', t', x') compatible with what Ralph imagines at t in w is such that x' never sees Ortcutt at t' in w' . Since there are centered worlds in which the center fails to see Ortcutt, this move would appear to avoid our worry.

But this proposal really just takes us back to the problematic view that we discussed at the beginning of §3, the singular proposition approach. So the problem for that view is equally well going to be a problem for the view just proposed. Suppose, for example, that Ralph imagines a scenario in which Ortcutt (relative to Q) is distinct from Ortcutt (relative to S). For example, suppose that Ralph is watching the gray-haired man on the beach, while at the same time holding a photograph of the man in the brown hat that he earlier saw behaving suspiciously. Then Ralph might have an imagining he could report by saying:

“I’m imagining a scenario in which I never saw that man [he points to the gray-haired man on the beach] nor this man [he points to the man in the photograph], and in which the two of them are having a fist-fight.”

The present proposal cannot make sense of this imagining, an imagining which is evidently coherent.

Note that the fist-fight case just mentioned is one that is problematic for both the descriptive account and the singular proposition account. This means that even if we adopt some kind of two-dimensionalism according to which attitudes have two kinds of content (descriptive and singular), we cannot evade the present problem.¹⁶ For the above case – a case which combines the problem of counterfactual attitudes with the main problem for the singular proposition approach – will be one in which both of the contents posited by the two-dimensionalist are empty.^{17,18}

5 Cheap haecceitism

On the centered worlds account of other-location, if Ralph has (at t in w) a *de re* attitude about Ortcutt (relative to Q), then each of Ralph’s attitudinal

¹⁶For the sort of view I have in mind, see, for example, Chalmers (2002).

¹⁷An anonymous referee suggests that the centered worlds theorist might resist this last argument by maintaining that Ralph’s imagining in the fist-fight case is not actually coherent, in spite of initial appearances. The referee suggests that the centered worlds theorist might describe the case as one in which Ralph is imagining a “disguisedly contradictory content.” This is a possible response, but it is not one that Lewis and his followers are likely to find congenial. As I understand the proposal, centered propositions are not believed or imagined *simpliciter*, but only believed or imagined under a disguise. But if Lewis regarded this sort of move as legitimate, then it seems mysterious that he should have gone down the descriptive route in the first place: why not say, of Ralph in Quine’s original case, that he believes the singular centered proposition that Ortcutt is a spy under one disguise, and that he believes the singular centered proposition that Ortcutt is not a spy under a different disguise? Clearly, the referee’s suggested proposal abandons one of Lewis’s ambitions, which is to represent intuitive distinctions between contents using no more than sets of possibilities. Maybe that ambition ought ultimately to be abandoned, but that more general issue exceeds the scope of the present paper.

¹⁸It is worth pointing out that, in the fist-fight case, one can characterize a *two-dimensional intension* (i.e. a function from centered worlds to sets of centered worlds) that is non-empty. For attempts to use this fact to solve the problem of counterfactual attitudes see Ninan (2008, Ch. 2) and Yanovich (2011).

alternatives contains an individual y who represents Ortcutt (relative to Q) there. The problem of counterfactual attitudes seems to arise because of a certain constraint the theory puts on what this individual y can be like. The constraint is this:

- a centered world (w', t', x') contains a representative of Ortcutt (relative to Q) only if there is an individual y' to whom x' bears Q uniquely at t' in w' ;
- if a centered world contains such an individual y' , then y' is the one who represents Ortcutt (relative to Q) in that centered world.

There is, in effect, no other way for an individual in one of Ralph's attitudinal alternatives to represent Ortcutt (relative to Q) as existing there. This leads straight to the problem of counterfactual attitudes, for if Ralph is to imagine (at t in w) anything about that man [we point at the brown-hatted man that Ralph is watching], he must imagine (at t in w) that he bears Q to that man. The constraint in effect requires that the individual who represents Ortcutt (relative to Q) in one of Ralph's attitudinal alternatives be *qualitatively similar* to Ortcutt himself in a particular way. More precisely, it requires that the pair (x', y') that represents the pair (Ralph, Ortcutt) in (w', t', x') be qualitatively similar to (Ralph, Ortcutt) in a certain way: x' must bear Q to y' at t' in w' just as Ralph bears Q to Ortcutt at t in w .

It seems that this requirement of qualitative similarity is the source of our trouble. So we should be able to avoid the trouble if we can solve the problem of 'de re representation' without requiring the relevant representatives to be qualitatively similar to what they represent in this way. To this end, it will be instructive to turn our attention briefly to a parallel discussion concerning the analysis of *metaphysical* modal claims.

On Lewis's well-known account of possible worlds, the domains of any two worlds are disjoint. Given that doctrine, *de re* modal claims like "Romney could have won the election" are not best understood as saying that there is a possible world in which Romney himself won the election. On Lewis's original version of counterpart theory (Lewis 1968), the above sentence would instead be understood as saying that there is a possible world w which contains a *counterpart* of Romney who won the election in w . Romney's counterparts in another world w are those things in w that are more similar to Romney than are any other things in w , and which are sufficiently similar to Romney. Multiply *de re* claims are handled in a similar fashion: "It could have been that Romney won and Obama lost" is true just in case there is a possible world w which contains a counterpart of Romney who won in w and a counterpart of Obama who lost in w .

In an early response to Lewis, Feldman (1971) argued that the following sentence might be true, but would come out false on the counterpart-theoretic analysis:

I could have been more similar to the way you in fact are than to the way I in fact am; and at the same time, you could have been

more similar to the way I in fact am than to the the way you in fact are.

The counterpart-theoretic analysis of this would be something like:

There is a possible world w and two individuals x and y in w such that: x is my counterpart in w ; y is your counterpart in w ; x is more similar to you than y is; and y is more similar to me than x is.

This seems to imply that y is more similar to me than x is. But then how could x be my counterpart in w ? Now there are probably a number of different ways for a defender of standard counterpart theory to respond to this objection, but I don't want to consider them here.¹⁹ Instead, I want to look at a solution to this problem proposed by Hazen (1979), a variant of which was later adopted by Lewis himself (cf. Lewis 1983a, 1986, §4.4). The view in question is sometimes called *cheap haecceitism*. What's interesting for us is that the cheap haecceitist proposal suggests a way of solving our problem of *de re* representation without entangling ourselves with the notion of qualitative similarity.²⁰

On Hazen's approach, possibilities – the things modal operators quantify over – are no longer identified with possible worlds. Instead, Hazen takes a possibility to be a pair (w, f) of a possible world w and a function f from the domain of the actual world into the domain of w . The idea is that f maps an individual x to his or her representative in the possibility given by (w, f) . (Hazen calls such functions *representative functions*.) So on Hazen's approach, "Romney could have won the election" is true just in case there is a possibility (w, f) such that $f(\text{Romney})$ wins the election in w . Similarly, "Romney could have won while Obama lost" is true just in case there is a possibility (w, f) such that $f(\text{Romney})$ wins in w and $f(\text{Obama})$ loses in w .

On Hazen's version of this view, a pair (w, f) of this sort is a genuine possibility only if, for any individual x , f maps x to something in w that is *of the same kind* as x . Lewis (1983a, 1986, §4.4), on the other hand, thinks that, at least in some contexts, there are no such constraints, and so regards every pair (w, f) of this sort as a genuine possibility.²¹ But on either view, Feldman's sentence should come out as true, since there are possibilities (w, f) such that

¹⁹For one such response, see (Lewis 1983b, 43).

²⁰Hazen's proposal is relevant to the problems for counterpart theory raised in Fara and Williamson (2005); see Dorr (2010), Bacon (ming), and Russell (ming). For additional discussion of cheap haecceitism, see Fara (2009), Kment (2012), and Stalnaker (2012).

²¹This is a bit misleading, since Lewis adopts a variation of Hazen's proposal according to which *de re* modal claims are analyzed by quantifying over sequences of compossible individuals that are accessible from a given sequence of actual individuals. But he thinks that, at least in some contexts, there are no constraints on this accessibility relation, in which case the resulting proposal is closely related to the proposal I attribute to Lewis in the text.

One difference between Lewis's view and Hazen's is that the former seems to allow an individual to have more than one representative in a given possibility. This is because the sequences of actual individuals that figure in Lewis's analysis of *de re* modal claims are allowed to contain repetitions. (But Hazen's approach could be modified by replacing representative functions with (possibly non-functional) relations.) For more discussion of the difference between Hazen-style proposals and Lewis's variant, see Dorr (2010, 16-17).

$f(\text{me})$ is more similar to you than to me, and such that, $f(\text{you})$ is more similar to me than to you. The reason for this is that the individual who represents me in (w, f) is no longer the individual in w who most closely resembles me – it is simply $f(\text{me})$, whoever that may be. The individual who represents me in a possibility is no longer determined by relations of qualitative similarity; it is, so to speak, simply stipulated that $f(\text{me})$ represents me in (w, f) , irrespective of how qualitatively similar (or dissimilar) the two of us are. For this reason, Hazen calls these (w, f) -pairs *stipulational worlds*.

6 Multi-centered worlds

Note a similarity here between stipulational worlds and centered worlds. When we are assessing a centered world (w', t', x') for compatibility with what an agent x believes or imagines (at time t in world w), the center x' represents the agent x at t' in w' . But x' represents x not in virtue of being identical to x nor in virtue of being qualitatively similar to x , but simply because this is the way the theory works—the theorist simply *stipulates* that the center of a centered world is the agent’s representative in that centered world. So here again we see the idea that the individual that represents someone in another possible situation is determined via theoretical fiat, rather than via some more robust relation, like identity or similarity.

But there is, of course, an important difference between stipulational worlds and centered worlds: the centered worlds theorist adopts this approach only for the agent’s *de se* representative; all other representatives must be qualitatively similar to the individual they represent in the manner described earlier. Now if the problem of counterfactual attitudes arises from this requirement of qualitative similarity, then perhaps we can avoid the problem by treating *de re* representation for attitudinal possibilities the way Hazen treats *de re* representation for metaphysical possibilities. But how do we do this?

Perhaps we can simply identify attitudinal possibilities with Hazen’s stipulational worlds rather than with centered worlds. But there are at least two obstacles to doing this. The first is that any individual x in the actual world has at most one representative $f(x)$ in any given stipulational world (w, f) – this follows from f ’s being a function. But because of cases of identity confusion – cases like the Ralph-Ortcutt case – we need to allow for the possibility that an attitudinal possibility contains more than one representative of a given individual in the agent’s world. In any of Ralph’s doxastic alternatives, there should be two distinct individuals both of whom correspond to Orcutt. Now we could avoid this problem if we replaced representative functions with ‘representative relations,’ i.e. relations which are possibly non-functional (cf. footnote 21). But this move wouldn’t avoid the second problem, which is that, when it comes to attitudinal possibilities, an individual in one possibility represents an individual in the agent’s world only relative to an acquaintance relation.

To appreciate this last point, think for a moment about how the centered worlds account handles the Ralph-Ortcutt case. In each of Ralph’s centered

belief worlds (w', t', x') , there is one individual y' who x' has seen in a brown hat at t' in w' , and there is a distinct individual z' who x' has seen on the beach at t' in w' . Both of these individuals represent or correspond to Ortcutt, but they do so relative to two different ways in which Ralph is acquainted with Ortcutt. The first individual, y' , represents Ortcutt relative to Q (the ‘see in a brown hat’ relation), while the second, z' , represents Ortcutt relative to S (the ‘see on the beach’ relation). So on the centered worlds theory, an individual in an attitudinal possibility does not represent an individual in the agent’s world full-stop; rather an individual in an attitudinal possibility represents an individual in the agent’s world only relative to an acquaintance relation.

This observation suggests a way we might adapt Hazen’s proposal to the case of attitudinal modalities.²² Instead of using functions from *individuals* to individuals (as Hazen does), we should instead use functions from *pairs of individuals and acquaintance relations* to individuals. We could then identify attitudinal possibilities with pairs of a possible world w and a function from (individual, acquaintance relation)-pairs to individuals in w . To see how this would work, suppose we gathered together all of the individual-acquaintance relation pairs (y, R) such that Ralph bears R to y at t in w . Let σ be the set of all of these pairs. The idea would be to represent Ralph’s attitudinal alternatives using pairs (w', f') consisting of a possible world w' and a function f' from σ into the domain of w' . The function f' would map each (y, R) in σ onto its representative in (w', f') . So $f'(y, R)$ will be the individual in (w', f') who represents y relative to R there. Since Ralph is acquainted (at t in w) with Ortcutt via Q and via S , both (Ortcutt, Q) and (Ortcutt, S) will be elements of our set σ . So in a possibility (w', f') , $f'(\text{Ortcutt}, Q)$ will represent Ortcutt relative to Q , while $f'(\text{Ortcutt}, S)$ will represent Ortcutt relative to S . Since f' can map each of those pairs – (Ortcutt, Q), (Ortcutt, S) – to different individuals in (w', f') , we can, like the centered worlds theorist, adequately represent Ralph’s belief state.

Let us develop these ideas more precisely, and then illustrate how they help to solve the problem of counterfactual attitudes. We begin by defining the notion of an *acquaintance set*:

Definition. An agent x ’s *acquaintance set* at time t in world w is the set of all pairs (y, R) such that x bears acquaintance relation R to individual y uniquely at t in w .

Definition. An *acquaintance set simpliciter* is the acquaintance set of some agent at some time in some world.

To distinguish the functions employed in our theory from Hazen’s representative functions, let us call the former *tagging functions*.

Definition. A *tagging function* is a (total) function from an acquaintance set into the domain of some possible world.

²²Note that nothing in what follows depends on Lewis’s idea that the domains of any two possible worlds are disjoint.

Where Hazen took metaphysical possibilities to be pairs of a possible world and a representative function, we will instead take the attitudinal alternatives of an agent x at a time t in world w to be triples of a possible world w' , a time t' and a tagging function f' , where f' is a function from x 's acquaintance set at t in w into the domain of w' . Let us call this general type of object a *multi-centered world*:

Definition. A *multi-centered world* is a triple (w, t, f) consisting of a possible world w , a time t , and a tagging function f whose range is included in the domain of w .

The term *multi-centered world* is appropriate because this approach extends the centered worlds treatment of *de se* representation to the general *de re* case. Each of the outputs of the tagging function f is a distinguished individual – a ‘center’ – of the multi-centered world (w, t, f) .²³

A multi-centered world (w', t', f') will be compatible with what an agent x imagines at time t in world w only if the domain of f' is x 's acquaintance set at t in w , and the individuals who are the outputs of f' have all the properties that the agent imagines that the corresponding inputs have. More precisely:

- (5) A multi-centered (w', t', f') is compatible with what an agent x imagines at t in w iff:
- the domain of f' is $\{(y_1, R_1), \dots, (y_n, R_n)\}$, x 's acquaintance set at t in w ,
 - each of the $f'(y_i, R_i)$ has, at t' in w' , all of the properties that x imagines, at t in w , that y_i has (relative to R_i), and
 - $f'(y_1, R_1), \dots, f'(y_n, R_n)$ stand, at t' in w' , in all of the relations that x imagines, at t in w , that y_1 (relative to R_1), ..., and y_n (relative to R_n) stand in.

Since any one-place property can be represented as an n -place relation ($n \geq 1$), we can drop the middle clause on the right-hand side of the biconditional.²⁴

As an alternative to the centered worlds account of (singly) *de re* imagining (cf. (3)), we offer the following:

- (6) An agent x imagines, at t in w , that y (relative to R) is F iff every multi-centered world (w', t', f') compatible with what x imagines (at t in w) is such that $f'(y, R)$ is F at t' in w' .²⁵

²³On the broad idea of having multiple centers (in some sense) in the representation of attitudinal alternatives, see Spohn (1996), Ninan (2008, Ch. 3), Ninan (2010), Torre (2010), and Chalmers (2011a,b).

²⁴Compare this gloss on the notion of compatibility to the one we gave for centered worlds in (1).

²⁵Since tagging functions total functions from acquaintance sets into worldly domains, it follows that any agent x who imagines something at time t in world w thereby imagines that y (relative to R) exists, where y is any individual to whom x bears some acquaintance relation R (at t in w). This is clearly false (and probably remains false even if substitute

Before illustrating how this account resolves the problem of counterfactual attitudes, I should point out how *de se* attitudes are treated on this account. We follow Lewis (1979) in taking *de se* attitudes to be *de re* attitudes where the *res* is the agent and the acquaintance relation is *I*, identity. So if (w', t', f') is compatible with what agent x believes or imagines at t in w , $f'(x, I)$ is x 's '*de se* representative' in (w', t', f') .

We can now see how this account avoids the problem of counterfactual attitudes. Recall this case:

Suppose Ralph is looking at the brown-hatted man before him, and has an imagining (at t in w) he could report by saying, "I'm imagining a situation in which I never laid eyes on that man [Ralph points at the brown-hatted man before him]. In my imaginary scenario, that man is a recluse who few people have ever seen."

Ralph is imagining (at t in w) that Ortcutt (relative to Q) is a recluse whom he has never seen. For Ralph to imagine this (at t in w) on our account is for all of the multi-centered worlds (w', t', f') compatible with what he imagines (at t in w) to be such that $f'(\text{Ortcutt}, Q)$ is a recluse at t' in w' who $f'(\text{Ralph}, I)$ has never seen at t' in w' . And there certainly are multi-centered worlds meeting this condition, since there are worlds w' and times t' such that some individual y' in w' is a recluse at t' in w' and some individual x' in w' has never seen y' at t' in w' . If (w', t', f') is such that $f'(\text{Ortcutt}, Q) = y'$ and $f'(\text{Ralph}, I) = x'$, then (w', t', f') is a multi-centered world meeting this condition.

That dispenses with the first problem involving counterfactual attitudes discussed earlier, the left-to-right problem. What about the second problem, the right-to-left problem? Recall the case:

Ralph bears Q to Ortcutt at t in w , but engages in a purely general imagining—he simply imagines that there is a unique man in a brown hat who he sees, and that this man is flying a kite. Here his imagining is intuitively not about Ortcutt.

This case is again no problem for the multi-centered account. For Ralph to imagine this (at t in w) is for all the multi-centered worlds (w', t', f') compatible with he imagines at t in w to be such that there is exactly one man y' in a brown hat who $f'(\text{Ralph}, I)$ is watching at t' in w' , and this man y' is flying a kite

"believes" for "imagines" in the preceding sentence). One way to avoid this result is to allow multi-centered worlds (w', t', f') where some of the outputs of f are objects that 'do not exist' at w' . What one thinks it means for an object to fail to exist at a world depends on one's background views about the metaphysics of modality. Those who think existence is contingent will interpret this to mean that we allow some outputs of f' to be individuals who do not exist at w' but are found in the domain of some other possible world. Those who think existence is necessary (e.g. Linsky and Zalta 1994, Williamson 1998) can interpret this to mean that we allow some outputs of f' to be individuals that are not concrete in w' . Alternatively, we could allow tagging functions to be partial functions, though this would require us to adjust the statement of (5), our gloss on the notion of what it is for a multi-centered world to be compatible with what an agent imagines. (Thanks to an anonymous referee for raising this issue.)

at t' in w' . In contrast, if Ralph were to imagine that he was watching Ortcutt (relative to Q) flying a kite, all of Ralph's imagination alternatives (w', t', f') would be such that $f'(\text{Ralph}, I)$ was watching $f'(\text{Ortcutt}, Q)$ flying a kite at t' in w' .²⁶

On the centered worlds account, if (w', t', x') is a centered world compatible with what Ralph imagines at t in w , then x' is the individual who represents Ralph (relative to I). If y' is the individual who represents Ortcutt (relative to Q) in (w', t', x') , the account demands that x' bears Q to y' , just as Ralph bears Q to Ortcutt. I have suggested that this requirement of qualitative similarity between (x', y') and (Ralph, Ortcutt) is what leads to the problem of counterfactual attitudes. Note how the multi-centered worlds account differs in this respect. Suppose (w', t', f') is a multi-centered world compatible with what Ralph imagines at t in w , and that $f'(\text{Ralph}, I) = x'$ and $f'(\text{Ortcutt}, Q) = y'$. Then x' represents Ralph (relative to I) and y' represents Ortcutt (relative to Q). But on this account, there is simply no constraint as to how (x', y') and (Ralph, Ortcutt) are related—there is simply no demand that those two pairs be qualitatively similar in any particular respect. It is this feature of the present account that enables it to avoid the problem of counterfactual attitudes.

Another way in which the present account differs from the centered worlds approach concerns what each account takes as its basic notion of belief or imagining, etc. (I will put the point in terms of belief, but the same goes, *mutatis mutandis*, for imagining and other attitudes.) For Lewis, the basic notion of belief is belief *de se*. This is the basic notion in two (related) senses: First, this is the notion of belief that one uses when assessing a potential doxastic alternative for compatibility with what an agent believes. Second, *de re* belief – believing something about someone relative to an acquaintance relation – is reduced to *de se* belief (cf. (2)). On the multi-centered approach, in contrast, the basic notion of belief is belief *de re*, or the notion of what an agent believes about an individual relative to an acquaintance relation (more generally: what an agent believes about a plurality of individuals-relative-to-acquaintance-relations). This is this notion of belief that we use when we characterize the notion of doxastic compatibility (cf. (5)). And *de se* belief is understood as a special case of *de re* belief – it is a *de re* belief where the *res* is the agent and the relation of acquaintance is I , identity. Thus, we have not tried to reduce the notion of *de re* belief to something more basic; unlike Lewis's account, the multi-centered account of *de re* belief is not reductive.

7 Limited accessibility

Although we have talked about the characterization of attitudinal alternatives on this approach, we have yet to say anything about what the *content* of an attitude is. Usually in a modal theory, the content of an attitude is a set of possibilities of some sort. So the natural thing to do on the multi-centered

²⁶The reader can verify that the present account has no trouble handling the fist-fight case discussed in §4.

approach is to take contents to be sets of multi-centered worlds, *multi-centered propositions*. To believe or imagine a multi-centered proposition p is for all of the multi-centered worlds compatible with what one believes or imagines to be contained in p .

But one consequence of moving from centered worlds to multi-centered worlds is that we lose the relativistic individuation of *de se* content that we discussed in §2. Recall the argument we discussed in favor of relative contents: Consider two agents, a and b , each of whom is in a belief state (at time t in world w) that she could express by saying, “I am being attacked by a bear.” These two agents will, other things being equal, be disposed to act in similar ways. So the contents of their respective beliefs should be the same. But since a ’s belief might be true while b ’s is false, the thing they believe needs to be something that can vary truth value across individuals. Lewis’s centered propositions have just this feature, and this was taken to be an argument for Lewis’s approach.

But consider what the multi-centered content of a ’s belief would be, if a believed (at t in w) that she was being chased by a bear. Individual a has a *de se* belief, which we construe as a *de re* belief about herself relative to I , the identity relation. So in each multi-centered world (w', t', f') compatible with what she believes (at t in w), her representative – $f'(a, I)$ – is being chased by a bear at t' in w' . So the content of her belief is the following set of multi-centered worlds:

$$(7) \{(w', t', f') : f'(a, I) \text{ is being chased by a bear at } t' \text{ in } w'\}$$

But the content of b ’s belief, on the other hand, is going to be something different. She has a *de re* belief about b , relative to I . So in each multi-centered world (w', t', f') compatible with what she believes (at t in w), her representative – $f'(b, I)$ – is being chased by a bear at t' in w' . So the content of her belief is:

$$(8) \{(w', t', f') : f'(b, I) \text{ is being chased by a bear at } t' \text{ in } w'\}$$

But since a is distinct from b , (7) and (8) are different sets. So a and b do not believe the same thing, despite the doxastic similarity between them.

In fact, it is not hard to see that, on the present approach, *de se* contents are what Perry (1979, 45) calls “propositions of limited accessibility.” For (7) is a multi-centered proposition that only *a* can believe, and (8) is a multi-centered proposition that only *b* can believe. To see this, suppose otherwise; suppose, that is, that an agent *x* distinct from *a* believes (7) at a time *t* in a world *w*. Let (w', t', f') be an arbitrary multi-centered world compatible with what *x* believes at *t* in *w*.²⁷ Since *x* believes (7) at *t* in *w*, it follows that (w', t', f') is in (7). But (w', t', f') is in (7) only if *f'* is defined for the pair (a, I) , i.e. only if (a, I) is in the domain of *f'*. Since (w', t', f') is compatible with what *x* believes at *t* in *w*, the domain of *f'* must be *x*'s acquaintance set at *t* in *w*. From these two facts it follows that (a, I) is in *x*'s acquaintance set at *t* in *w*, and from this it follows that *x* bears *I* to *a* at *t* in *w*. Since *I* is the relation of identity, this means that *x* is identical to *a*. But this contradicts our assumption that *x* is distinct from *a*. So that assumption must be false; only *a* can believe (7). Parallel reasoning shows that only *b* can believe (8).

It is worth pointing out that, on this account, it is *de se* contents specifically whose accessibility is limited in this way. Two agents *a* and *b* can believe the same *de re* and *de dicto* contents in the appropriate circumstances. I mention this because one might be tempted to object to the present account as follows:

Consider two agents *a* and *b* at time *t* in world *w*. We know that *a* and *b* don't have the same acquaintance set, since *a*'s acquaintance set contains (a, I) , whereas *b*'s does not (since *b* is not identical to *a*). Now consider an arbitrary multi-centered world (w', t', f') that is compatible with what *a* believes at *t* in *w*. Given the account of multi-centered compatibility (cf. (5)), *f'*'s domain is *a*'s acquaintance set. Since *a*'s acquaintance set differs from *b*'s, it follows that *f'*'s domain is not *b*'s acquaintance set. So it follows that (w', t', f') is not compatible with what *b* believes at *t* in *w* (again, by (5)). Since (w', t', f') was an arbitrary multi-centered world compatible with what *a* believes at *t* in *w*, this result holds for all multi-centered worlds compatible with what *a* believes at *t* in *w*, i.e. it follows that no multi-centered world compatible with what *a* believes at *t* in *w* is compatible with what *b* believes at *t* in *w*. But then *a* and *b* cannot believe any of the same multi-centered propositions. It is one thing for *de se* contents to be ‘private’—perhaps we can live with this. But it is quite another thing for all contents to be private; surely that is absurd.²⁸

I agree that this is absurd, but I deny that it is a consequence of the present account.

²⁷If no multi-centered worlds are compatible with what *x* believes, then it is vacuously true that *x* believes (7), even if *x* is distinct from *a*. We leave this degenerate case aside; the interesting point is that *a* is the only agent with a coherent, non-trivial belief state that can believe (7).

²⁸Thanks to Dirk Kindermann and an anonymous referee for raising this concern.

It is true that no multi-centered world compatible with a believes at t in w is compatible with what b believes at t in w . But it doesn't follow from this that there is no multi-centered proposition that both a and b believe. Consider an example. Suppose a and b both see Ortcutt in a brown hat, and both bear relation Q to Ortcutt at t in w . While observing this brown-hatted man, both a and b come to believe (at t in w) that this man (relative to Q) is a spy. I claim that both a and b believe (at t in w) the following multi-centered proposition:

$$r = \{(w', t', f') : f'(\text{Ortcutt}, Q) \text{ is a spy at } t' \text{ in } w'\}$$

If the objection is correct this should be impossible. But why should it be impossible? All of the multi-centered worlds compatible with what a believes are in r ; and all of the multi-centered worlds compatible with what b believes are in r ; yet no multi-centered world compatible with a believes is compatible with what b believes. All of the dogs are in the house; all of the cats are in the house; yet no dog is a cat.

The key point is that two multi-centered worlds (w'', t'', f'') and (w''', t''', f''') might both be in r even while the domain of f'' differs from the domain of f''' . The set r collects up all those multi-centered worlds (w', t', f') that have the following property: the individual that f' maps $(\text{Ortcutt}, Q)$ to is a spy at t' in w' . The only requirement this places on the domain of f' is that it be defined for $(\text{Ortcutt}, Q)$. So even if the domain of f'' differs from the domain of f''' , both functions might be defined for $(\text{Ortcutt}, Q)$, in which case both (w'', t'', f'') and (w''', t''', f''') might be elements of r .²⁹

8 Recovering relativism

So on this approach, the contents of *de re* and *de dicto* attitudes are not in general propositions of limited accessibility. But the contents of *de se* attitudes are. Is even this more restricted result problematic? It does, of course, conflict with Lewis's motivation for positing relative contents. But the proponent of the multi-centered approach has two possible responses to Lewis. The first is to capitulate. This is possible, because as I shall presently demonstrate, there is a way of amending our approach so that it yields relative *de se* contents. The second response is to resist Lewis's objection. After all, Lewis's position is not uncontroversial. Frege (1956) famously held the view that the content of an agent's *de se* attitude was something that could only be believed by the agent herself, a view defended by some contemporary philosophers, such as Evans (1981). I shall first demonstrate how one might capitulate, and then explain how one might resist. I close with a tentative discussion of which response (if either) ought to be preferred.

²⁹Note that on the centered worlds approach, if I believe that I am short, and you believe that you are *not* short, the set of centered worlds compatible with what I believe will be disjoint from the set of centered worlds compatible with what you believe. For all of my centered belief worlds will be centered on short individuals, while all of yours will be centered on non-short individuals. That doesn't mean there are no centered propositions that we both believe. (Similar remarks apply to the possible worlds theory.)

In our development of the multi-centered theory, we have been taking the ‘*de re* representation relation’ to be a relation between individuals in a possibility and individual-acquaintance relation pairs (y, R) , where the agent bears R *uniquely* to y . Let me emphasize that last point: all of the acquaintance relations we’ve been talking about are ones that the agent bears uniquely to some individual (at the time and world in question).

Now suppose we have a multi-centered world (w', t', f') , where the domain of f' is some agent x ’s acquaintance set at t in w . And suppose we do the following: we take each element (y, R) of x ’s acquaintance set and ‘throw out’ the individual component of the pair. What we’re left with then is a set of acquaintance relations: that set of all acquaintance relations R such that x bears R uniquely to someone at t in w . We can make a parallel adjustment to the function f' : let the domain of f' be the set of acquaintance relations $\{R_1, \dots, R_n\}$ obtained in this manner. I submit to you that, despite having been transformed in this way, the triple (w', t', f') can still be regarded as representing a potential attitudinal alternative for x at t in w . To see this, let “ $R^{w,t,x}$ ” abbreviate “the individual to whom x bears R uniquely at t in w ”. Then we can regard each of the $f'(R_i)$ as representing $R_i^{w,t,x}$ relative to R_i .

To put it loosely: since we can always ‘recover’ an individual from an acquaintance relation (together with the agent, time, and world), we can throw the individuals out of the agent’s acquaintance set, and everything will work just as before. The main difference between a theory formulated along these lines and the theory developed in the previous section of the paper concerns the individuation of content. In particular, the version of the theory in which individuals are purged from acquaintance sets yields a relativistic individuation of *de se* content.

Instead of taking an acquaintance set to be a set of individual-relation pairs, we will now take an acquaintance set to be a set of acquaintance relations:

Definition. An agent x ’s acquaintance set at t in w is the set of all acquaintance relations R such that there is someone to whom x bears R uniquely at t in w .

We leave the other definitions as they were: a tagging function is still a function from acquaintance sets to individuals in the domain of some world, and a multi-centered world is a triple (w, t, f) , where f is a tagging function. But since acquaintance sets are now sets of acquaintance relations, tagging functions are now functions from acquaintance relations to individuals.

Compatibility is then characterized as follows (I use the state of belief here, rather than the state of imagining):

- (9) A multi-centered (w', t', f') is compatible with what an agent x believes at t in w iff:
- the domain of f' is $\{R_1, \dots, R_n\}$, x ’s acquaintance set at t in w , and
 - $f'(R_1), \dots, f'(R_n)$ stand, at t' in w' , in all of the relations that x believes, at t in w , that $R_1^{w,t,x}$ (relative to R_1), ..., $R_n^{w,t,x}$ (relative to R_n) stand in.

The account of (singly) *de re* belief on this version of the proposal goes like this:

(10) An agent x believes, at t in w , that y (relative to R) is F iff:

(i) $R^{w,t,x} = y$,

(ii) every multi-centered world (w', t', f') compatible with what x believes (at t in w) is such that $f'(R)$ is F at t' in w' .

It is straightforward to verify that this version of the multi-centered approach resolves the problem of counterfactual attitudes in much the same manner as the first version does.

The content of an attitude is a set of multi-centered worlds, a multi-centered proposition. But now with our new understanding of what acquaintance sets, tagging functions, and multi-centered worlds are, the content of a *de se* attitude will be relativistic in the way centered propositions are. To see this, suppose again that, at t in w , a believes *de se* that she is being chased by a bear, and b believes *de se* that she is being chased by a bear. Then every multi-centered world (w', t', f') compatible with what a believes (at t in w) will be such that $f'(I)$ is being chased by a bear at t' in w' (recall that I is the identity relation). So, at t in w , a believes the following multi-centered proposition:

(11) $\{(w', t', f') : f'(I) \text{ is being chased by a bear at } t' \text{ in } w'\}$

But everything is exactly the same for b : every multi-centered world (w', t', f') compatible with what she believes (at t in w) will be such that $f'(I)$ is being chased by a bear at t' in w' , and so b too will count as believing (11).

But in spite of believing the same thing, a 's belief might be true, even while b 's is false. To see how the account allows for this possibility, we need to say something about what it is for a belief to be true on this account:

Let $\{R_1, \dots, R_n\}$ be x 's acquaintance set at t in w , and

let $f^{w,t,x} = \{(R_1, R_1^{w,t,x}), \dots, (R_n, R_n^{w,t,x})\}$

If x believes a multi-centered proposition p at t in w , then x 's belief is true iff $(w, t, f^{w,t,x}) \in p$.

Suppose a is being chased by a bear at t in w , but that b is not. And suppose that a 's acquaintance set at t in w is $\{(I, I^{w,t,a}), (R_2, R_2^{w,t,a}), \dots, (R_n, R_n^{w,t,a})\}$, where I is the identity relation. Then, by the above account, her belief is true iff:

$(w, t, f^{w,t,a}) \in (11)$,

where $f^{w,t,a} = \{(I, I^{w,t,a}), (R_2, R_2^{w,t,a}), \dots, (R_n, R_n^{w,t,a})\}$

So a 's belief is true iff $f^{w,t,a}(I)$ is being chased by a bear at t in w . Note that $f^{w,t,a}(I) = I^{w,t,a}$, and that $I^{w,t,a}$ is the individual to whom a bears I at t in w , i.e. a . So a 's belief is true iff a is being chased by a bear at t in w . Since a is being chased by a bear at t in w , her belief is true. Parallel reasoning shows that, when b believes (11) at t in w , her belief is true iff b is being chased by a

bear at t in w ; since we are supposing that b is not being chased by a bear at t in w , her belief is false. So a and b believe the same thing, even though a 's belief is true, while b 's is false.

So we now have two versions of the multi-centered account: one which, like Fregean accounts, takes *de se* contents to be propositions of limited accessibility; and one which, like the centered worlds account, takes them to be relative propositions, things that vary in truth value between individuals. Which account should we adopt? I will not try to settle this issue conclusively here, but I will try to give the reader a sense of the relevant considerations. The proponent of relativistic content can, of course, point to the 'bear attack' argument in defense of her position. The first premise in that argument is that two agents who are in the same *de se* belief state – for example, two agents who are both in a belief state that they could express by saying, "I am being attacked by a bear" – are disposed to behave in similar ways (other things being equal). The second premise is that, in virtue of the fact reported by the first premise, we ought to say that the two agents believe the same thing.

Let us agree that the role *de se* attitudes play in the explanation of action requires us to accept that agents in similar *de se* states will, other things being equal, be disposed to act in the same way. *One* way to capture this similarity would be to say that they believe the same thing. But another way to capture it would be to say that, although such agents don't believe the same thing, they share a common property, the property of being in the same *de se* belief state. It is this property that is relevant for explaining why they are disposed to behave similarly. For example, someone who wanted to defend the 'limited accessibility' version of the multi-centered approach could point out that, in our bear attack example, both a and b have the property of being an x such that x believes:

$$\{(w', t', f') : f'(x, I) \text{ is being chased by a bear at } t' \text{ in } w'\}$$

The fact that they share this property might then be relevant for explaining why they are disposed to behave similarly, perhaps in conjunction with other properties of this sort that they share (e.g. the property of believing that one can avert a bear attack by playing dead, the property of wanting to avert a bear attack, etc.). Unless more is said, it seems that considerations from the explanation of action do not tell decisively in favor of the relativistic version of the multi-centered approach.

A further difference between the two accounts, which might be relevant to how we choose between them, concerns *externalism*. The first version of the multi-centered approach is externalist in a way in which the second view is not. To illustrate this, consider two individuals, Jones and Smith, on opposite ends of the globe. Jones bears some acquaintance relation R uniquely to a particular individual—Jane, say. Smith finds himself in an epistemic situation which, 'from the inside', is exactly like the situation Jones finds himself in. And he too bears R uniquely to some individual—Susan (where Susan is distinct from Jane). Now suppose that Jones comes to believe that Jane (relative to R) has brown hair; Smith similarly comes to believe that Susan (relative to R) has brown hair.

According to the first version of the multi-centered approach, the content of Jones's belief is (12), while the content of Smith's is (13):

(12) $\{(w', t', f' : f'(Jane, R) \text{ has brown hair at } t' \text{ in } w')\}$

(13) $\{(w', t', f' : f'(Susan, R) \text{ has brown hair at } t' \text{ in } w')\}$

Since Susan is distinct from Jane, these are distinct contents, so what Jones believes differs from what Smith believes. But on the second version of the multi-centered approach, the two men believe the same thing:

(14) $\{(w', t', f' : f'(R) \text{ has brown hair at } t' \text{ in } w')\}$

In this respect, the second version of the account is again more similar to the centered worlds approach, which would also say that Jones and Smith believe the same thing.

References

- Bacon, A. (Forthcoming). Representing counterparts. *Australasian Journal of Logic*.
- Chalmers, D. J. (2002). The components of content. In D. Chalmers, editor, *Philosophy of Mind: Classical and Contemporary Readings*. Oxford University Press, New York.
- Chalmers, D. J. (2011a). Frege's puzzle and the objects of credence. *Mind*, **120**(479), 587–635.
- Chalmers, D. J. (2011b). The nature of epistemic space. In A. Egan and B. Weatherson, editors, *Epistemic Modality*, pages 60–107. Oxford University Press, Oxford.
- Dorr, C. (2010). How to be a modal realist. Unpublished manuscript, Oxford University.
- Egan, A. (2010). Disputing about taste. In T. Warfield and R. Feldman, editors, *Disagreement*. Oxford University Press, Oxford.
- Evans, G. (1981). Understanding demonstratives. In H. Parret and J. Bouveresse, editors, *Meaning and Understanding*. Walter de Gruyter, Berlin. Reprinted in Evans 1985, 291–321. Page references are to the 1985 reprint.
- Evans, G. (1985). *Collected Papers*. Clarendon Press, Oxford.
- Fara, D. G. (2009). Dear haecceitism. *Erkenntnis*, **70**(3), 285–297.
- Fara, M. and Williamson, T. (2005). Counterparts and actuality. *Mind*, **114**(453), 1–30.
- Feldman, F. (1971). Counterparts. *Journal of Philosophy*, **68**(13), 406–409.
- Frege, G. (1918/1956). The thought: A logical inquiry. *Mind*, **65**(259), 287–311.
- Hazen, A. (1979). Counterpart-theoretic semantics for modal logic. *The Journal of Philosophy*, **76**(6), 319–338.

- Hintikka, J. (1962). *Knowledge and Belief*. Cornell University Press, Ithaca, N.Y.
- Kaplan, D. (1968). Quantifying in. *Synthese*, **19**(1-2), 178–214.
- Kment, B. (2012). Individuals, supervenience, and fundamentality. *Philosophical Review*, **121**(4), 573–609.
- Kripke, S. (1980). *Naming and Necessity*. Blackwell, Oxford.
- Lewis, D. K. (1968). Counterpart theory and quantified modal logic. In Lewis (1983b), pages 113–126.
- Lewis, D. K. (1979). Attitudes *de dicto* and *de se*. *Philosophical Review*, **88**(4), 513–543. Reprinted in Lewis 1983b, 133–159. Page references are to the 1983 reprint.
- Lewis, D. K. (1983a). Individuation by acquaintance and by stipulation. *Philosophical Review*, **92**(1), 3–32. Reprinted in Lewis 1999, 373–402. Page references are to the 1999 reprint.
- Lewis, D. K. (1983b). *Philosophical Papers, Volume 1*. Oxford University Press, New York.
- Lewis, D. K. (1986). *On the Plurality of Worlds*. Blackwell, Oxford.
- Lewis, D. K. (1995). Reduction of mind. In S. Guttenplan, editor, *A Companion to Philosophy of Mind*, pages 291–324. Basil Blackwell, Oxford.
- Lewis, D. K. (1999). *Papers in Epistemology and Metaphysics*. Cambridge University Press, Cambridge.
- Linsky, B. and Zalta, E. (1994). In defense of the simplest quantified modal logic. *Philosophical Perspectives*, **8**, 431–458.
- Moss, S. (2012). Updating as communication. *Philosophy and Phenomenological Research*, **85**(2), 225–248.
- Ninan, D. (2008). *Imagination, Content, and the Self*. Ph.D. thesis, MIT.
- Ninan, D. (2010). *De se* attitudes: Ascription and communication. *Philosophy Compass*, **5**(7), 551–567.
- Ninan, D. (2012). Counterfactual attitudes and multi-centered worlds. *Semantics and Pragmatics*, **5**(5), 1–57.
- Perry, J. (1977). Frege on demonstratives. *Philosophical Review*, **86**(4), 474 – 497. Reprinted in Perry 1993, 3–32. Page references are to the 1993 reprint.
- Perry, J. (1979). The problem of the essential indexical. *Noûs*, **13**(1), 3–21. Reprinted in Perry 1993, 33–52. Page references are to the 1993 reprint.
- Perry, J. (1993). *The Problem of the Essential Indexical and Other Essays*. Oxford University Press, New York.
- Perry, J. (2001). *Reference and Reflexivity*. CSLI Publications, Stanford.
- Quine, W. (1956). Quantifiers and propositional attitudes. *Journal of Philosophy*, **53**(5), 177–187. Reprinted in Quine 1976, 185–196.
- Quine, W. (1976). *The Ways of Paradox and Other Essays*. Harvard University Press, Cambridge, MA.

- Russell, J. S. (Forthcoming). Actuality for counterpart theorists. *Mind*.
- Spohn, W. (1996). On the objects of belief. In C. Stein and M. Textor, editors, *Intentional Phenomena in Context. Papers from the 14th Hamburg Colloquium on Cognitive Science.*, volume 55, pages 117–141.
- Stalnaker, R. (1984). *Inquiry*. MIT Press, Cambridge, MA.
- Stalnaker, R. (1999). *Context and Content*. Oxford University Press, New York.
- Stalnaker, R. (2008). *Our Knowledge of the Internal World*. Oxford University Press, Oxford.
- Stalnaker, R. (2012). *Mere Possibilities*. Princeton University Press, Princeton, NJ.
- Torre, S. (2010). Centered assertion. *Philosophical Studies*, **150**(1), 97–114.
- Williamson, T. (1998). Bare possibilities. *Erkenntnis*, **48**, 257–273.
- Yanovich, I. (2011). The problem of counterfactual *de re* attitudes. *Semantics and Linguistic Theory*, **21**, 56–75.