

Who's Your Ideal Listener?

Ethan Nowak and Eliot Michaelson

Abstract

It is increasingly common for philosophers to rely on the notion of an idealized listener when explaining how the semantic values of context-sensitive expressions are determined. Some have identified the semantic values of such expressions, as used on particular occasions, with whatever an appropriately-idealized listener would take them to be. Others have argued that, for something to count as the semantic value, an appropriately-idealized listener should be able to recover it. Our aim here is to explore the range of ways that such idealization might be worked out, and then to argue that none of these results in a very plausible theory. We conclude by reflecting on what this negative result reveals about the nature of meaning and responsibility.

1 Introduction

Suppose I point to the Vice-Chancellor of my university and say:

- (1) That's the bastard who's trying to cut my pension.

What makes the Vice-Chancellor the referent of my utterance of 'that'? At a first pass, it might seem like enough to say that the Vice-Chancellor is the first object intersected by the ray drawn through my outstretched finger.¹ On reflection, however, any number of rays might be drawn through my finger—at a sufficient

This work is entirely collaborative. Pre-print version; please cite the published version.

¹See Quine (1968) and McGinn (1981).

distance, these will intersect different objects. Even if we limit our attention to a single ray, it is not easy to see how to guarantee that the Vice-Chancellor turn out to be the referent of my demonstrative, given that any ray that intersects him is likely to intersect his finely-tailored jacket or his hand-finished shirt first.

In light of these problems, some philosophers claim that a speaker's intentions should be invoked to supplement her gestures.² Instead of arguing about rays, the story goes, we should simply ask which object the speaker intended to target with her gesture. This move, however, leaves us with little to say about the fact that (1) can be used perfectly naturally without any gesture at all, or about cases in which a gesture fails completely to indicate the object the speaker intends it to target.³

Many philosophers deal with such possibilities by jettisoning gestures entirely and saying that speaker's intentions fix the reference of demonstratives directly.⁴ This, however, raises a different worry: surely we cannot use bare demonstratives like 'this' or 'that' to refer to whatever we like whenever we like! Pointing at a stuffed parrot in the Pitt River Museum but thinking about a cat you knew years ago in Thailand, you cannot now utter the words 'that cat' and succeed in getting this demonstrative to pick out the cat.

The most popular response in the face of this sort of difficulty has been to rein in the speaker's intentions, allowing them to fix reference only when they meet particular constraints. Bach (1992) requires that speakers' intentions must exhibit a kind of Gricean reflexivity. Others have argued that this is insufficient, and appeal instead

²Cf. Kaplan (1978) and Reimer (1992).

³See Reimer (1992).

⁴Cf. Bertollet (1980), Schiffer (1981), Davies (1982), and Kaplan (1989). See also Russell (1948) for a nearby view based on attention rather than intention.

to certain facts about the listener. Stokke (2010), for example, thinks reference succeeds only when the actual listener recovers the speaker's intended referent. Since this proposal entails that reference fails whenever the listener happens not to be paying attention, many have preferred to appeal to *idealized* listeners. King (2014), to mention a paradigmatic example, argues that a token demonstrative like 'this' or 'that' will refer to some object x only if (i) the speaker intends to refer to x with her use of the relevant demonstrative, and (ii) a suitably-idealized listener would be in a position to recognize that x was the intended referent.⁵

Interestingly, similar appeals to idealized listeners have surfaced elsewhere in the literature on metasemantics, in work on the reference of indexicals like 'I' and 'now'. In order to resolve a number of hard cases that arise when such terms are recorded, Romdenh-Romluc (2002, 2006) proposes that indexical tokens refer to whatever a suitably-idealized listener would take them to refer to. So when 'I' is played back on my answering machine, we say that it refers to me, since that is who a suitably-idealized listener will take the term to refer to once she recognizes my voice and finds that it matches the person she intended to call.⁶

Our primary aim here will be to argue that the reactions of a suitably-idealized listener cannot be used to fix the referent of a singular term, nor can they plausibly serve as part of a larger theory by providing a necessary condition for reference. We will make the case for this claim by showing that, no matter how the notion of an 'ideal listener' is spelled out, there are examples of successful reference which that sort of ideal listener will be unable to recover. This task will be made tractable by our demonstrating that, at root, there are just two basic idealization strategies.

⁵See Evans (1973), Wettstein (1984), Strawson (2000), and Neale (2004) for similar suggestions.

⁶For a criticism of this proposal, see Weatherson (2002). For a summary of the issues involved in answering-machine cases, see Cohen and Michaelson (2013).

This will allow us to offer a fairly intuitive characterization of why neither can capture all of the cases.

Having shown where ideal listener theories come up short, we close by considering what this entails for the notion of ‘doing enough’ to ensure communication. While there is no space here to explore the connections in the detail they deserve, we take the issues we have identified in the case of singular terms to reflect broader debates about context sensitivity, and indeed, possibly about linguistic meaning generally.⁷ By showing how the idea of ‘responsible speech’ might be decoupled from the project of semantics per se, we hope to set the stage for more productive dialogue going forward.

2 Ideal listeners are hard to find

Although the notion of an ideal listener lies at the heart of the theories of meaning we are interested in, it has not typically been spelled out in great detail. Here are the three most extensive elaborations of which we are aware:

[T]he coordination account claims that the value of a given use of a demonstrative is that object that meets the following two conditions: 1) the speaker intends it to be the value; and 2) a competent, attentive, reasonable hearer would take it to be the the object that the speaker intends to be the value. We should understand condition 2 in such a way that the idealized hearer knows the common ground of the conversation in Stalnaker’s sense at the time the sentence containing the

⁷Thanks to an anonymous referee for reminding us that the argument we describe about demonstratives really echoes even older and more general debates about meaning between Griceans, for example, and relevance theorists.

demonstrative was produced. (King 2014, pg. 226)

[W]e should separate (i) *who or what A intended to refer to by an expression X on a given occasion*, and (ii) *who or what a rational, reasonably well-informed interpreter in B's shoes thinks A intended to refer to by X on that occasion*. In cases where (i) = (ii), we can talk freely about *what the speaker referred to*. (Neale 2004, pg. 80, emphases in original)⁸

[W]e should hold that indexical reference must be fixed by a competent and attentive audience—Ac. Note that Ac is not just any competent and attentive bystander who happens to hear the utterance. As mentioned above, U will sometimes exploit the beliefs, desires, history, interests, etc. of her audience as a means of indicating the reference-determining context, hence Ac is the audience that it is reasonable to take U to be addressing. (Romdenh-Romluc 2006, pg. 274)⁹

Each of these formulations involves the idea that an ideal listener must be attentive and linguistically competent. King requires, furthermore, that an ideal listener be reasonable, and Neale that she be rational. While Romdenh-Romluc makes no explicit stipulation to this effect, we imagine that the listeners she has in mind will be guided by the same basic principles of rationality that are standardly taken to apply to the participants in a conversation (by e.g. Davidson, Grice, and others).

The most substantial divergences between the views described by King, Neale, and

⁸A caveat: as Neale (2016) later clarifies, 'freely talking about what the speaker referred to' is not an endorsement of the standard notion of semantic reference. Indeed, Neale is working within a framework where there is really *only* speaker's reference (though, even in this later work, he is happy to call one type of speaker's reference 'expression reference [relative to a context]'). We bracket further consideration of Neale's framework in what follows, since his ideas regarding successful reference should be of interest regardless of whether one accepts the rest of that framework.

⁹Romdenh-Romluc (2002) also contains some very similar ideas, though never put quite as directly as they are in the quoted passage.

Romdenh-Romluc concern the ways in which their idealized listeners are modeled on participants in the concrete speech situations in which the utterance takes place. On King's presentation, the ideal listener is stipulated to know the common ground of the conversation, in the sense of Stalnaker (2002). Since the common ground is determined by the epistemic states of the agents who participate, King's ideal listener will inherit some of the general features of the state of mind of the actual speaker and listener, but is not otherwise required to be psychologically or epistemically related to any actual agent. The formulations offered by Neale and Romdenh-Romluc suggest what we take to be a range of different possible positions, which vary depending on how directly the ideal listener is linked to the actual listener (and indeed, to the actual speaker, as well).

Consider the idea at the heart of Neale's proposal, according to which a speaker refers to an object with a demonstrative just when she intends to do so and her intention is one that a 'reasonably well-informed interpreter in [the listener's] shoes' would recognize. While we can imagine reasonable people disagreeing about what kind of agent should count as reasonably well-informed in a context—disagreeing with regard to features that might result in different predictions about reference—for now, we want to focus our attention on the enormous variety of ways in which an agent might count as being in another agent's shoes. On a very loose construal, you might count as occupying someone else's shoes when you occupy the same position in space that they do, and are thus presented with the same scene from the same perspective. On such a construal, Neale's proposal sounds a lot like King's—the test of reference is the test of whether, if you took a random English speaker of ordinary intelligence, made them stand in the place the actual hearer is standing, and look in the direction the actual hearer is looking, they would understand

which object the speaker had in mind.

On the other hand, Neale's formulation could equally straightforwardly apply to positions that require much more substantial similarities between the ideal listener and the person who actually hears the demonstrative utterance. Although we hesitate to place much weight on folk intuitions about what it means to 'be in someone else's shoes', it strikes us that, as the expression is typically used, a person only counts as occupying the position when they know what the other person knows, or have lived experiences similar to those of the other person, or enjoy the same perceptual affordances and evaluative responses as the other person, and so on. (If you see a child struggling to find a way to move a heavy object, you don't think 'If I were in his shoes, I'd just pick it up!') So we think it is just as easy to read Neale as offering a treatment of referential terms according to which the ideal listener is one who is granted not just access to the position in space occupied by the actual listener, but access furthermore to her memories, perceptual experiences, dispositions, interests, evaluations, and so on. On the strongest variant of this view, the ideal listener might diverge from the actual listener only in that she is guaranteed to be attentive whereas the actual listener is not—and perhaps guaranteed to avoid errors in her demonstrative reasoning.

On the stronger construals of the relationship between the ideal listener and the actual listener, Neale's position sounds more like the position we are inclined to read into Romdenh-Romluc's work than it does like King's. Romdenh-Romluc emphasizes that, on her view, it will not do to evaluate the semantic values of singular referential terms from the perspective of a random bystander who happens to be dropped into the scene. Her ideal listener is clearly supposed to be a version of the actual listener, one equipped to pick up on the ways in which the speaker can

‘exploit the beliefs, desires, history, interests, etc. of her audience as a means of indicating the reference-determining context’ (loc. cit.).

Romdenh-Romluc’s description of the ideal listener, however, brings out another dimension according to which both her view and Neale’s appear to admit of a range of variants. In addition to her explicit commitment to the idea that the ideal listener should be broadly and deeply similar to the actual listener, she endorses another theme that appears to pull in a different direction. After observing that speakers very often exploit highly specific information about their interlocutors in order to refer, she concludes that the ideal listener should be the listener that it would be ‘reasonable to take [the speaker] to be addressing’ (loc. cit.).

We can see two ways of understanding this idea. On the first of these, the ideal listener would be defined by taking the actual listener—with all of her capacities, interests, and history—and then guaranteeing that she is being attentive and barring her from making errors in her reasoning. This is much like the strong version of Neale’s view; effectively, it idealizes the listener in a way that should seem reasonable to an outside observer.

On the other salient reading, which we take to be the one Romdenh-Romluc more likely had in mind, the ideal listener is a listener endowed with the capabilities, interests, and history that the speaker has good reason to attribute to her, regardless of whether she actually possesses these various qualities. This would allow the speaker to take advantage of more idiosyncratic features of the listener’s memory and perspective, or at least what the speaker has good reason to take these to be these to be, in order to establish reference.

To summarize, the range of views suggested by King, Neale, and Romdenh-Romluc

might usefully be grouped into two broad categories. On what we will call ‘non-specific idealization’ views (like King’s view, as well as the weaker readings of Neale), the ideal listener is supposed to represent the person in the street. The basic idea behind these views is that, for an object to count as the referent of a particular utterance of a demonstrative, it must be whatever object a random, normal, attentive listener who had access to some general information about the conversation (like the common ground, or the location in which it takes place) would identify as the referent. On what we will call ‘specific idealization’ views (like the natural readings of Romdenh-Romluc and the stronger readings of Neale), in contrast, the ideal listener is characterized using more specific information about the actual parties to the conversation. In the following two sections, we will argue that, regardless of how exactly the details are filled in, both the non-specific and the specific approach to ideal listeners make the wrong predictions about a wide range of cases involving demonstratives.

The basic dilemma, as we see it, is this: if you appeal to what a fairly normal listener, idealized primarily with respect to attentiveness and reasonableness, would recover in a context, then you are forced to classify cases where the speaker takes advantage of the idiosyncrasies of her listener in order to coordinate on a referent as cases of reference failure. If, on the other hand, you idealize the listener in such a way as to maintain those idiosyncrasies, then you risk having to classify a range of cases where the speaker has intuitively ‘done enough’ to secure reference, despite a de facto lack of coordination, as cases of reference failure.

3 ‘Non-specific’ idealizations will not do

As Romdenh-Romluc points out in the passage quoted above, speakers commonly leverage highly specific facts about their audiences when using demonstratives. We take this to constitute a crippling objection to ideal listener theories that involve non-specific idealizations. Consider an example from David Robson, writing for the BBC:¹⁰

Superhuman Vision

As Concetta Antico took her pupils to the park for an art lesson, she would often question them about the many shades she saw flashing before her eyes. I’d say, ‘Look at the light on the water—can you see the pink shimmering across that rock? Can you see the red on the edge of that leaf there?’ The students would all nod in agreement. It was only years later that she realised they were just too polite to tell the truth: the colours she saw so vividly were invisible to them.

Today, she knows that this is a symptom of a condition known as ‘tetrachromacy’. Thanks to a variation in a gene that influences the development of their retinas, people like Antico can see colours invisible to most of us. Consider a pebble pathway. What appears dull grey to you or me shines like a jeweller’s display to Antico. ‘The little stones jump out at me with oranges, yellows, greens, blues and pinks,’ she says. ‘I’m kind of shocked when I realise what other people aren’t seeing.’

Imagine that the rock that shimmers pink for Antico is just one among many rocks

¹⁰Available at: <http://www.bbc.com/future/story/20140905-the-women-with-super-human-vision>.

in a river, the rest of which look uniformly—to both her and her students—the way ordinary rocks do to a trichromat observer. From her perspective, then, the pink rock stands out clearly.

Consider what counts as the common ground in the conversation after Antico's initial questions. She asked her students whether they can see the rock with the pink shimmer in the river, and they nodded in response. Antico's question presupposes that each listener can see a rock with a pink shimmer on it, and that presupposition was accommodated when no one objected to it by saying 'huh?' Since accommodated presuppositions are standardly assumed to be added to the common ground unless they are explicitly challenged, we take this to mean that, on any plausible construal of the common ground, it is now mutually presupposed *that each interlocutor can see a rock in the river with a pink shimmer*.¹¹

With this background in place, consider what happens when Antico follows up on her earlier questions with:

- (2) Everyone sketch a picture of that rock with the pink shimmer.

This is a serious challenge for King's version of the non-specific idealization theory. For King, recall, an object is the referent of a demonstrative just in case it is both what the speaker intends as the referent and what a competent, attentive, and reasonable listener who knows the common ground of a conversation would take to be the referent. In this case, Antico has a referent in mind when she says 'that rock with the pink shimmer', but it is one that a competent, attentive, and reasonable,

¹¹Remember that a proposition need not be mutually believed in order to enter the common ground of a conversation. So, the fact that only Antico really believes that everyone can see the pink rock makes no difference. Compare Stalnaker (2002): 'It is common ground that X in a group if all members accept (for the purpose of the conversation) that X, and all believe that all accept that X, and all believe that all believe that all accept that X, etc.' (pg. 716).

though otherwise normal (trichromatic) listener will have no hope of recovering. In other words, we have a real-life case where a reasonable, competent, and attentive listener who knows the common ground of a certain conversation will be unable to identify the object that her interlocutor intends to pick out with her utterance of a complex demonstrative—not because she lacks some bit of knowledge or competence, but because she lacks some basic perceptual capacity.

We think this is a bad result. Intuitively, to appropriate a favored locution of King's, Antico has 'done what she can' to refer. We find it hard to imagine a plausible standard by which she could be considered an irresponsible demonstrative-user, and can thus see no obvious reason for attributing reference failure. Still, since we imagine that some philosophers will not share our intuitions on this point, it is worth making clear that the case can be modified slightly so as to make the problem substantially worse.

Consider a variation on the scenario described above which makes all of Antico's students tetrachromats themselves, and where Antico also knows this (she has, suppose, given them all a color-perception test beforehand). In such a case, the students will all share the kind of perceptual experience that Antico has when looking at the shimmering pink rock, and thus will have no trouble identifying the rock she has in mind. If, in the presence of this exceptional group, Antico should utter (2), we submit that the result would be as clear a case of successful reference as ever there is. A particular object stands out to the speaker in a striking way. She has reason to expect that the object will stand out in the same way for her audience, and it in fact does. She intends to say something about the object, and her audience takes her to have done so.

Nevertheless, King's account predicts that Antico's utterance of 'that rock with the

pink shimmer' will again fail to refer to the rock that both she and her interlocutors visually identify in this way. Although her interlocutors in fact understand which object she has in mind, her intended referent is not one that an ordinary (trichromatic) listener who knows the common ground of the conversation would be able to identify, no matter how attentive that listener were to be.¹² So King predicts that her utterance of this complex demonstrative refers to nothing. We find this result intolerable, and we take it to show that King's proposal is wrong.

We think the tetrachromat example is a particularly nice one because it is not a philosopher's fiction, and because we know that tetrachromatic vision is not normal in humans on anyone's characterization of normal. But lest we be seen as resting our case on what might seem like a marginal example—even if it is a real one—we should emphasize that similar examples can be generated using a variety of pedestrian resources. Consider just a few illustrations:

Listener's Preferences

I have invited you to a cake shop to celebrate your birthday. Without your knowing, I've done some sleuthing to discover that your favorite cake is Boston Cream Pie. So neither your preference nor my knowledge thereof is a part of the common ground.

We enter the shop, and are greeted by a large number of delicious-looking offerings. I say: 'Look—that's your favorite, isn't it?' intending to refer to the Boston Creme Pie, but not making any gesture or doing

¹²One might try insisting that it is tetrachromats who should count as 'normal' relative to this conversational context. To flesh out this sort of suggestion, however, we would need to specify which factors determine what counts as a 'normal' audience member relative to a context. The natural suggestion would be to idealize the listener in line with how the speaker reasonably takes her to be; in other words, this suggestion turns out to be a variant of the 'specific idealization' view, which we will consider in detail in the next section. Thanks to [ommitted] for discussion.

anything else to make clear which cake I have in mind. Nevertheless, you know which cake I am talking about, because you know what your favorite cake is.

We take this to be as plausible a case of referential success as any. I leverage a piece of knowledge I have about you to make my referential intention public. You recognize my intention and correctly identify the object I have in mind. But a linguistically competent, reasonable, attentive individual who knew the common ground of the conversation would not be able to do this. So, for King, my utterance of ‘that’ fails to refer.

Shared Memories

Riding the cable car in La Grave, we are presented with an extremely complex high-alpine panorama. Without gesturing, I say: ‘Look—that’s the line we skied last year, isn’t it?’ intending to refer to a particular route that weaves through cliff bands and a series of narrow couloirs. You know which line I am talking about, because you were there (and you have the skill required to read the terrain—more on that later).

Surely the proposition that we skied a particular visible route does not have to be part of the common ground (we can easily imagine that the listener did not recall skiing in the area in question until being reminded of it) for this conversation to proceed. And even if that is part of the common ground, it is unlikely to enable any random listener, even one skilled in reading terrain, to pick out the right route. So, we take it that someone who knew the common ground would not thereby be in a position to identify which object I have in mind. But my interlocutor here will easily recover the object I have in mind, and we take that to be a good indication

that reference succeeds.¹³

Specialized Vocabulary

A first responder appears at the scene of a wilderness medical emergency. A second person happens by, sees the first attending to the victim, and sees an open backpack filled with first-aid supplies. Neither one knows anything about the other, so the common ground is minimal. The first responder says: ‘Hand me that SAM splint.’ Fortunately, the second person knows what a SAM splint is and complies.

Most people will not know what a SAM splint is. This means the token demonstrative, on King’s view, will have no referent. We think that, in this case, the demonstrative pretty clearly does have a referent. So we take this case to tell against King’s view.¹⁴

Specialized Skills

We work in a microbiology lab. I have just come across a particularly interesting sample. I say: ‘Wow, here’s a wild one. Check this thing out!’ You have a look through the microscope, and immediately recognize the particular object that has attracted my attention.

If you are an experienced viewer of objects through this kind of microscope, you are in possession of a specific set of skills that not everyone is in possession of. While

¹³One might object that the relevant common ground to know isn’t the one at the start of the utterance, but after any presuppositions are accommodated. Although this conflicts with King’s claim that the relevant common ground is the one ‘at the time the sentence containing the demonstrative was produced’ (loc. cit.), we take the suggestion to be worth exploring. However, given that accommodation will depend on highly specific features of the listener (e.g. her memories), we will treat this as a variant of the ‘specific idealization’ view.

¹⁴King might say that his requirement that an ideal listener be linguistically competent rules this out as a counterexample, on the grounds that you would have to know what a SAM splint is to be a competent user of the expression. But we don’t think this is a plausible requirement. In general, we see little reason to accept that a linguistically competent interpreter must be competent with the specific predicates from which any given complex demonstrative is formed.

it may be plausible to think that it is part of the common ground here that we are both microbiologists, we do not think that entails that anyone who had access to the common ground would possess the requisite skill itself. Even people who know a lot about microbiology—including professional microbiologists—might not themselves know how to interpret what they see in the microscope (they might work with a different sort of microscope, for instance). So, anyone who thinks that this is a plausible case of successful reference should also take it to be a case that King's account makes the wrong predictions about. For King, the fact that an ordinary person who knew the common ground of the conversation would not be able to identify the object the speaker has in mind means that the speaker's utterance of 'that thing' has no referent.

We think each of the examples presented so far, from *Superhuman Vision* through to *Specialized Skills*, are equally effective against the version of Neale's view on which reference succeeds only when someone standing in the listener's position would be able to track the referent. Standing where a person with tetrachromatic vision is standing and looking at the same scene from her position is clearly not enough to allow you to see what she sees. The referential cues that will be available to a tetrachromat will not be available to a normal observer in her position in space. Of course, the same goes for standing in the position occupied by a person who knows where another person skied last year, or who can count on their having a certain approach to microscopy or whatever.

Our examples also undermine one of the interpretations we have offered for Romd-enh-Romluc (an interpretation which, we noted, was not likely one she herself would endorse). On that interpretation, a speaker's utterance of a demonstrative refers to whatever object an outside observer would take an idealized version of

the listener to be most likely to recover as the referent. In *Superhuman Vision*, we take it that this procedure will yield the wrong result. An ordinary outside observer will look at the scene in question and think: ‘What pink rock? Where?’ Then she will conclude, on the basis of this, that no idealized version of the actual listener is going to be able to identify any particular rock as the referent. On reasonable assumptions, the same goes for most of our other examples; a normal person off the street, for example, will have no idea what a SAM splint is. Nor will she have any reason to think that the listener will know what a SAM splint is. In fairness, this version of Romdenh-Romluc’s theory does look better positioned than either King’s or Neale’s to deal with *Listener’s Preferences*: for, even though the outside observer herself will have no idea what the listener’s favorite cake is, she will have every reason to believe that the listener herself does know, and hence that an idealized version of that listener will be able to identify that cake as the referent.

4 ‘Specific idealizations’ won’t work either

The cases we have looked at so far all underscore the significance of a point that Romdenh-Romluc makes in the passage cited at the start of section 2: frequently, specific information about the people involved in a conversation matters where the question of reference is concerned. King’s view, and the versions of Neale’s that go in for what we might call the ‘a random person in the street would understand’ criterion for reference, miss this point. We think this disqualifies them as options.

This by itself, however, does not show that ideal listeners have no place in a theory of reference. The version of Neale’s view on which the ideal listener is one who

shares the audience's perceptual experiences, her memories, and so on, will survive the counterexamples we have presented so far. The same goes for the version of Romdenh-Romluc's view on which a speaker refers to an object just in case the listener she believes herself (with good reason) to be addressing is a listener whom she could (reasonably) expect to understand her; ditto for the closely related ways of spelling out Neale's view. There is even a strand in King's thinking that he could fall back on in light of the examples we have presented so far; although we think the position King explicitly describes is clearly wrong, he frequently claims that his aim is to throw light on what counts as having 'done enough' to qualify responsible demonstrative user.¹⁵ Perhaps 'doing enough' could instead be understood as per the surviving versions of Neale or Romdenh-Romluc's views.

So, should we accept an ideal listener theory based on what we might call, in contrast to the previously-described mechanism of non-specific idealization, 'specific idealization'? That is, should we allow ideal listeners access not just to the Stalnakerian common ground of a conversation, but to a richer range of epistemic and perceptual states had by the audience—or reasonably believed by the speaker to be had by the audience—as well?

We think not. We are not at all convinced that an idealized version of the actual listener—one who shares her perceptual experiences, memories, etc., and who is then idealized however one likes—really needs to be able to recover the speaker's intended referent in order for reference to succeed. Likewise, we doubt that you really need to have a reasonable expectation that your audience will be able to recover the object you have in mind in order to succeed in referring to that object.

¹⁵See King (2013, pg. 290) and King (2014, pp. 225, 229).

To bring out the first point, consider *Specialized Vocabulary* again. On Romdenh-Romluc's view, demonstrative reference succeeds only if an idealized version of the listener that the speaker takes herself to be addressing—a fully competent and attentive version of that listener—would identify the SAM splint as the referent of the speaker's utterance of 'that SAM splint'. But since the speaker will take it to be unlikely for her listener to know what a SAM splint is, this means that reference in this case will fail on Romdenh-Romluc's view. Note that this holds regardless of whether, as in our original case, the listener happens (by chance) to know what a SAM splint is, and hence *can* identify the referent. In other words, Romdenh-Romluc's view predicts that reference fails even when the speaker and listener manage to coordinate on an object, non-accidentally, in virtue of the speaker's utterance.

On a version of the specific idealization view like Neale's, on the other hand, successful reference requires coordination on a referent with an idealized version of the actual listener—regardless of what the speaker takes the listener to be like. In a case like *Specialized Vocabulary*, we take this to mean that reference will succeed if the actual listener happens to know what a SAM splint is, and will fail otherwise. This strikes us as an odd result, however, for it means that the truth or falsity of a speaker's utterance can depend on whether her actual listener happens to possess some highly specialized bit of vocabulary.

Our other cases from above can all be modified to bring out this sort of problem for the specific idealization view. Suppose, in a variation on *Shared Memories*, that the speaker doubts her listener will remember having skied the relevant line last year. Still, the speaker might hope that seeing the panorama will jog her listener's memory. If they succeed in coordinating, do we really want to call this a case of

reference failure? Or take *Specialized Skills*. Suppose I know that you are the wrong kind of biologist to be able to recognize the specimen through this microscope, but I plan to verbally guide you once you look through the lens. We think it is much more natural to say that in such a case I am planning on bringing you to see what my use of ‘that’ referred to than it is to say that I referred to nothing, but was aiming to lay the foundation for some possibly successful future reference.

Things get worse from here. Consider a case of what we might call ‘Hail Mary’ reference, where a speaker rightly believes that the environment will make communication all but impossible, but where she tries anyway—as there is nothing else for her to do.

Après Ski

A ski patroller sees a group of clueless skiers sitting down to have lunch under a massive overhanging serac. Recognizing the serious danger, she has to act—but she can tell from the rented airbag backpacks, unbuckled helmets, and sparkling outerwear that these are British tourists, likely monolingual, and she is a monolingual French speaker. Nevertheless, she shouts in French: ‘That is a dangerous area—get out of there now!’ She is right: these are monolingual British tourists and they have no idea what she is trying to convey.

In this case, the speaker has reason to believe that her audience will not understand her. Nevertheless, should a catastrophe ensue, prompting a subsequent investigation, we take there to be a very real sense in which our frustrated patroller can claim: ‘I said it was dangerous, but they wouldn’t budge’. We think, in other words, that there is a very real sense in which one can ‘do enough’ to refer without that guaranteeing successful coordination, even by an idealized version of the actual

audience.¹⁶

Moreover, we think there is a very real sense in which one can *know* all this and still be doing enough to refer. On this sense of ‘doing enough’, which we take to be the one most plausibly applicable to the question of reference, what ‘doing enough’ means is something like ‘doing all you could do’. Sometimes doing all you can do will not be enough, even when your audience is being attentive, reasonable, etc.; sometimes, all of this will even be clear to you ahead of time.¹⁷

So far we have focused on cases involving agents who intuitively succeed in referring to a certain object despite their rightly having no expectation that their interlocutors will recognize this (and thus, despite the fact that the proposed idealizations would not resolve the issue). But we can distort the speaker’s understanding of her situation in the other direction too. Consider:

Missing the Picture

I point behind myself to where my picture of David Cameron usually hangs and utter ‘That is a picture of someone who is likely to go down in history as Britain’s worst ever Prime Minister.’ In fact, there is a picture of Trump hanging behind me. I know this, having seen the switch. I also have good (though misleading) reasons to believe that you know it too. So I intend for you to use your knowledge to work out my intended referent. I am mistaken, however; you are unaware

¹⁶We grant that the ski patroller may not have succeeded in *warning* the tourists. As Austin (1962) claims, it might be that a successful warning requires uptake. Reference, we take it, comes at the level of locution, and is plausibly not so dependent.

¹⁷Note that we do not dispute that referring to *x* requires *trying* or even *intending* to get one’s audience to recognize *x* as the referent. We take it that one can intend to ϕ even when one harbors doubts regarding whether one will succeed; and certainly, one can try even when one expects to fail (cf. Ludwig 1992).

of the switch and have no idea what I'm talking about.¹⁸

We take it that, even given these beliefs, the speaker's use of 'that' fails to refer to the picture of Cameron as opposed to the picture of Trump. But the version of the specific idealization view that relies on the speaker's (reasonable) beliefs about the listener predicts the opposite. If we idealize as per this version of the theory, the idealized listener should recognize the speaker's intention to refer to the picture of Cameron.

A minor modification to the case will make it problematic for versions of the 'specific idealization' view that ignore the speaker's reasonable beliefs about the listener and focus instead on making the actual listener attentive and error-free in her reasoning. We simply stipulate that the listener knows about the switch, and can thus recover the speaker's intended referent. We have no doubt that, in this case, the *speaker* manages to refer to her picture of Cameron. But we doubt that the speaker's utterance of 'that' so refers. Rather, we take this to be a paradigm case of speaker's reference parting ways from semantic reference. You just can't point directly at a picture of Trump and refer to an absent picture of Cameron by saying 'that'—even if you manage to convey what you're thinking to your listener. This version of the specific idealization view predicts the opposite, however.

We take it, then, that the specific idealization strategy fails. Nevertheless, the theory points towards something we think is often overlooked but importantly right: in theorizing about meaning and communication, the specifics of the people involved matter. So while we reject this particular way of taking the lesson to heart, we hope to gesture towards a more promising way of doing so.

¹⁸This is, quite obviously, a variant on Kaplan (1978)'s famous Carnap-Agnew case (p. 239).

5 Conclusion

While the foregoing has been largely negative, we close with some more positive suggestions. We take it that one of the basic motivations for ideal listener-type theories of reference—and indeed, listener-based theories of meaning more generally—is to make reference depend on the speaker’s intuitively ‘doing enough’ to secure coordination. ‘Non-specific’ versions of the theory ask that she do this relative to a generic listener; ‘specific’ versions of the theory ask that she do this with respect to the actual listener, or at least with respect to what she reasonably takes that listener to be like. We argued that neither of these uniformly captures our judgments about reference—perhaps, at least in part, because neither uniformly captures what we mean by ‘doing enough’ to ensure coordination.

Neither of us is convinced that a speaker is actually required to ‘do enough’ in order for demonstrative reference to succeed. If what we called Hail Mary reference is possible, then we are unsure whether this notion is even well-defined in every case. What’s more, there are cases involving evasive testimony and similar where the speaker’s goal seems to be precisely to preclude the listener from recovering the intended referent—even though it is also important to the speaker to be testifying truthfully.¹⁹ If we take these cases at face value, they suggest that ‘doing enough’ to secure communicative success is not actually necessary to fix the values of referential terms like demonstratives.

Nonetheless, we certainly do think that ‘having done enough’ is an important evaluative status that we naturally apply to speakers when they use referential terms (and, indeed, a range of other context-sensitive expressions and possibly linguistic

¹⁹See Weatherson (2002) and [omitted] for further discussion of cases like these.

expressions generally). As such, we take this to be a worthwhile target for philosophical theorizing. To explain this status, it may again be tempting to appeal to idealized listeners. But we suspect it will be more fruitful to replace those idealizations with a claim along the following lines: what it means to ‘do enough’ will vary depending on how exactly we are interested in holding the speaker to account. In other words, it might be that what counts as ‘doing enough’ will depend on the sort of overall evaluative project we are pursuing when we ask whether the standard has been met on this or that occasion.

So, for instance, what counts as ‘doing enough’ to communicate with a demonstrative in a context looks like it will vary quite a bit depending on the payoff structure of communicative failure versus success, and the costs of doing more to ensure successful communication. For example, if we are writing a contract, we are likely to build quite a bit of descriptive material into any complex demonstrative we use, in order to avoid any possibility of confusion. If a contract writer fails to do this, we may judge her rather harshly for not having ‘done enough’; we might even go so far as to claim that the law should act as though reference here failed, regardless of whether or not the relevant interpreters manage to track the right objects. On the other hand, when the benefits of success are high but the costs of failure low, as in our Hail Mary cases, it seems as though we may be willing to employ extremely low standards when judging whether someone has ‘done enough’. At the very least, we seem loathe to judge speakers harshly in such cases, even if they have not done what would have been required to get their point across.

If ‘doing enough’ to ensure communication with a demonstrative really is this context sensitive, then we can imagine a friend of the ideal listener theory attempting to rehabilitate the theory by saying something similar: what it means to be ideal is

relative to a certain set of interests we have in asking about what a term refers to. This would entail a fairly radical rethinking of the metasemantic project—making it far more interest-relative than most have taken it to be to this point, more like the sort of deeply contextualist metasemantics that Travis (1996) proposes—but it might also serve to make the view more plausible. To be clear, we do not mean to be endorsing such a view. Nonetheless, if one is drawn to the sorts of motivations that have led many to think that ideal listeners will have a significant role to play in a theory of meaning, then we take this to be an idea worth considering.²⁰

References

Austin, J. L. (1962). *How to Do Things with Words*. Harvard University Press, Cambridge.

Bach, K. (1992). Intentions and demonstrations. *Analysis*, 52:140–146.

Bertolet, R. (1980). Demonstrations and intentions. *Philosophical Studies*, 38:75–78.

Cohen, J. and Michaelson, E. (2013). Indexicality and the answering machine paradox. *Philosophy Compass*, 8(6):580–592.

Davies, M. (1982). Individuation and the semantics of demonstratives. *Journal of Philosophical Logic*, 11:287–310.

Evans, G. (1973). The causal theory of names. *Proceedings of the Aristotelian Society*, 47:187–225.

Kaplan, D. (1978). Dthat. In French, P., Uehling, T., and Wettstein, H., editors, *Con-*

²⁰For helpful discussion and feedback on earlier drafts of this essay, many thanks to Jonathan Cohen, Sam Cumming, Nat Hansen, and Andrew Peet. Thanks also to two anonymous referees for this journal.

- temporary Perspectives in the Philosophy of Language*, pages 383–400. University of Minnesota Press, Minneapolis.
- Kaplan, D. (1989). Afterthoughts. In Almog, J., Perry, J., and Wettstein, H., editors, *Themes from Kaplan*, pages 565–614. Oxford University Press.
- King, J. C. (2013). Supplementives, the coordination account, and conflicting intentions. *Philosophical Perspectives*, 27(1):288–311.
- King, J. C. (2014). Speaker intentions in context. *Noûs*, 48(2):219–237.
- Ludwig, K. (1992). Impossible doings. *Philosophical Studies*, 65:257–281.
- McGinn, C. (1981). The mechanism of reference. *Synthese*, 49:157–186.
- Neale, S. (2004). This, that and the other. In Reimer, M. and Bezuidenhout, A., editors, *Descriptions and Beyond*, pages 68–182. Oxford University Press, Oxford.
- Neale, S. (2016). Silent reference. In Ostertag, G., editor, *Meaning and Other Things: Essays on Stephen Schiffer*, pages 229–342. Oxford University Press, Oxford.
- Quine, W. V. (1968). Ontological relativity. *The Journal of Philosophy*, 65(7):185–212.
- Reimer, M. (1992). Three views of demonstrative reference. *Synthese*, 93:373–402.
- Romdenh-Romluc, K. (2002). Now the French are invading England! *Analysis*, 62(1):34–41.
- Romdenh-Romluc, K. (2006). ‘I’. *Philosophical Studies*, 128:257–283.
- Russell, B. (1948). *Human Knowledge*. George Allen & Unwin, London.
- Schiffer, S. (1981). Indexicals and the theory of reference. *Synthese*, 49:43–100.
- Stalnaker, R. (2002). Common ground. *Linguistics and Philosophy*, 25:701–721.

- Stokke, A. (2010). Intention-sensitive semantics. *Synthese*, 175:383–404.
- Strawson, P. F. (2000). Singular direct reference: Intended reference and actual reference. In *Entity and Identity, and Other Essays*, pages 92–99. Clarendon Press, Oxford.
- Travis, C. (1996). Meaning's role in truth. *Mind*, 100:451–466.
- Weatherson, B. (2002). Misleading indexicals. *Analysis*, 62(4):308–310.
- Wettstein, H. K. (1984). How to bridge the gap between meaning and reference. *Synthese*, 58(1):63–84.