

Are there epistemic conditions necessary for demonstrative thought?

Michael Barkasi

October 23, 2019

Abstract

Starting with Gareth Evans, there's an important tradition of theorizing about perception-based demonstrative thought which assigns necessary epistemic conditions to it. Its core idea is that demonstrative reference in thought is grounded in information links, understood as links which carry reliable information about their targets and which a subject exploits for demonstrative reference by tokening the mental files fed by these links. Perception, on these views, is not fundamental to perception-based demonstrative thought but is only the information link exploited in these cases. Evans himself assigns a further epistemic condition (knowledge of a target's location in public space), while more recently Imogen Dickie has expanded the reliability requirement into a more complex account centered around justification. In this paper I synthesize three central proponents of this approach (Evans, Recanati, and Dickie) and show that the epistemic conditions they place on perception-based demonstrative thought are not actually required. My argument gives two examples in which there is perceptual contact with an object but this perceptual contact fails to do the epistemic work in question. The first case is stimulus-incorporating dream experiences, the second involves multimodal binding failures. I argue that this perceptual contact still affords demonstrative thought in these cases.

1 Introduction

What enables the perception-based demonstrative selection of referents for thought? One approach, starting with Gareth Evans, holds that all demonstrative selection is enabled by information links, understood as links which carry reliable information about their targets and which a subject exploits to select a referent for thought by tokening the mental files fed by these links. These links are *epistemically robust* (as I'll say) both because of their reliability and because the information they provide is usable by a thinker in virtue of how the links feed mental files, which are subject-level constituents of thought.

In this paper I will focus on three versions of this (as I'll call it) epistemic approach: Evans' seminal *The Varieties of Reference* (1982), Recanati's *Mental Files* (2012), and Dickie's *Fixing Reference* (2015a). Evans and Dickie add further layers of epistemic requirements on demonstrative selection: Evans says that it requires knowledge of the target's location in public space, while Dickie says that the links exploited for selection must provide justification for beliefs. I will show, by developing intuitions on two examples, that perception-based demonstrative selection does not require epistemically robust information links, nor do Evans' and Dickie's further requirements hold either.

The problem can be brought out with a rough-and-ready thought experiment. Focus attention on something you can see or hear clearly in your vicinity. Now imagine that your perceptual experience was introspectively identical to what it is now, but was epistemically inert: it provided no reliable information about the thing, no justification for beliefs, or knowledge of the thing's spatial location. Assuming the case still involves successful perception, wouldn't this experience enable the demonstrative selection of the thing you perceive for thought? Could you not still attend to what you perceive and think about it in a way which would be exhaustively expressed by uttering 'that'?

It will be tempting for my opponent to deny the assumption that such a case, as described, would still involve genuine perceptual contact. They may try to assimilate this case, on the basis of the epistemic failure, to cases of straight-up perceptual failure like hallucination and, thus, deny that your experience enables demonstrative selection. The two examples I provide below are designed to preserve perceptual contact while some epistemic feature (important to the view under discussion) is lost. In both cases I will sug-

gest that this perceptual contact suffices for demonstrative selection. The common problem among epistemic approaches is that they mistake incidental epistemic features of typical perceptual contact (that it provides reliable information, justification, or knowledge of spatial location) for essential conditions of demonstrative selection.

Two caveats: First, I am not denying that thought (including demonstrative thought) has epistemic requirements of some minimal sort. For example, I am not denying that thinking a thought requires some general capacity of the subject to deploy that supposed thought within their reasoning. I am also not denying that thoughts themselves (either thought vehicles or content) necessarily have certain features, such as fitness for use in justifications or decomposability into redeployable constituent parts. My concern is specifically with the epistemic requirements on demonstrative selection imposed by the epistemic approach.

Second, this is not a paper about how reference is fixed within perceptual states themselves. I do not attempt to give any explanation of how our perceptual states come to be about, or present to us, the particular stimuli we experience in them.¹ I will simply take for granted that our perceptual states do put us in contact with the world; they do present to us particular stimuli, or have particular content involving those stimuli. In large part the topic of my paper is whether these perceptual states (or the broader causal connections in which they figure) must, in order to enable demonstrative selection of referents for thought, satisfy additional epistemic requirements imposed by the epistemic approach.

2 Background

Perception-based demonstrative thoughts can be introduced via example (Smithies 2011, 7). Imagine that in front of you is an item. You look at it and think a thought which would be natural to express by asking “what’s *that*?” while pointing at it. A neutral gloss is that perception-based demonstrative thoughts are thoughts about something perceived, reference to which is enabled through a one-off episodic capacity afforded by that perception (Martin 2002, 178–81), and which are natural to express with demonstrative terms like ‘that’ (Dickie 2015b, 833). Although it’s often assumed (e.g., Bach 2010,

¹Important work on this question is done by Campbell (2002), Matthen (2005), Raftopoulos and Müller (2006), and Pylyshyn (2007).

55; Levine 2010, 179; Recanati 2012, 57–67) that these thoughts involve indexical mental representations (so-called *mental demonstratives*), others reject this assumption (e.g., Millikan 2012). I’ll drop the ‘perception-based’ and hereafter call them ‘demonstrative thoughts’, which is standard practice (e.g., Campbell 2002, 114; Pylyshyn 2007, 18; Dickie 2010, 213; Levine 2010, 169). Demonstrative thoughts can take the form of both standing propositional attitudes (e.g., beliefs) and mental action tokens (e.g., discrete instances of judging, inferring, wondering, etc), although my examples will be of the latter form.² It’s widely taken that demonstrative thoughts ground other derivative forms of mental and linguistic representation (Evans 1982, 181; Siegel 2006, 429; Pylyshyn 2007, 18–20; Levine 2010, 169; Smithies 2019, 66), and if so they are obviously very important.

For any demonstrative thought, there must be some facts which fix its referent (Pylyshyn 2007, 3–8, 23; Levine 2010, 172; Dickie 2015a, 115). When you think “what’s *that*?” while looking at the unknown item, what makes the unknown item (and not something else) the thing about which you’re wondering?³ It’s generally held (sometimes with argument) that demonstrative thoughts are not descriptive but are relational (e.g., Evans 1982, 173; Campbell 2002, 7; Bach 2010, 55; Jeshion 2010, 134; Recanati 2012, 12, 37; Dickie 2015a, 115; see also Dickie 2015b). Since relational thoughts are (by definition) thoughts which attribute a property to some particular fixed by a relation between that particular and the thinker, the answer to this question will involve some story about a relation between you and the unknown item.⁴ This relation enables the *selection* of the particular thing at the other end for thought and, when exploited, fixes this thing as the referent of the resulting demonstrative thought.

What is the relation which enables demonstrative selection? A tempting approach (see Barkasi 2019) is to appeal to the perceptual states involved in the thought: e.g., it is your focused attention to the item as it’s presented in your visual experience which fixes it as the referent of your thought. Being

²This division isn’t meant to be exhaustive or rigorous.

³For simplicity I’ll refer to thoughts by the utterances used to express them.

⁴It’s generally further assumed that relational thoughts are *singular*. What this means is that the subject component of the thought’s content contains the actual referent itself, or perhaps an object-dependent Fregean sense referring to that referent, and not a description which the referent satisfies. Less technically, what it means is that in thinking the thought the subject manages to think of the referent “directly” (in some sense), i.e. not by first conceptualizing the referent under some description which it satisfies.

the experienced target of your focused attention is a way for an item to be related to you, and it seems that you exploit this relation in demonstrative thought. It's fair to read John Campbell as holding this view (see [Campbell 2002, 2004](#); see also [Clark 2006](#)). Just about everyone agrees that attention has *something* to do with how the referents of demonstrative thoughts get fixed (e.g., [Evans 1982, 175](#); [Levine 2010, 179](#); [Dickie 2011, 294](#); [Matthen 2012](#), section 6). It's also fairly common to hold that, specifically, attention functions in cognitive tasks as a means of selection (e.g., [Pylyshyn 2007](#); [Levine 2010](#); [Carrasco 2011](#); [Wu 2011](#)); so the idea that the referents of demonstrative thought are fixed via selection by attention is not a stretch, although it has been criticized (e.g., [Levine 2010, 178](#)).

Still, the epistemic approach introduces an epistemic element to the story. It accepts that demonstrative thought is relational and that the reference-fixing relation is perceptual in some way: e.g., that it somehow involves attention to the object, perceptual experience of the object, or the operation of sensory systems. But on this approach, the perceptual nature of this relation is merely incidental. Imogen Dickie, in an early articulation of her view, expresses the idea well when she says that “perceptual attention enables perceptual demonstrative reference *because* it secures canonical *justification* for basic uses of a perceptual demonstrative that is luck-eliminating with respect to the attended object” ([2011, 295](#), emphasis added). Even without the details the idea should be clear: what's really important for enabling the selection of items (i.e., fixing of referents) in demonstrative thought is securing some epistemic condition—e.g., justification.⁵

Different versions of the epistemic approach will identify different epistemic conditions as essentially tied up in demonstrative selection. These accounts make some epistemic condition(s) *necessary* for demonstrative selection. My claim is that the epistemic approach is wrong: the epistemic conditions it imposes on demonstrative selection do not hold. Perceptual links do not need to satisfy, or enable the subject to satisfy, the additional epistemic requirements imposed by versions of the epistemic approach in order

⁵Even Campbell, whom I cited above as the prime example of an attention-based view, adds an epistemic element to his account. He assumes that selection of targets for demonstrative thoughts “is what causes and justifies the use of particular procedures to verify and find the implications” of those thoughts ([2002, 25](#)), and thereby that whatever fills this role of selection must also fill this accompanying epistemic role. But unlike Dickie, Campbell is less clear on whether this role of causing-and-justifying is incidental or essential to the selection process.

to enable demonstrative selection. These links do not enable demonstrative selection *because* they satisfy these further epistemic requirements.

3 Evans and Knowledge of Spatial Location

Although the epistemic approach might be traced back to Bertrand Russell and his notion of acquaintance (see [Russell 1912](#)), Evans' posthumous *The Varieties of Reference* ([1982](#)) provides the first well-developed version. In it Evans ([1982](#), 121) observes that “Our particular-thoughts are very often based upon information which we have about the world.” Such a thought is “governed by a conception of [its object] which is the result ... of a belief about how the world is which the subject has because he has received information (or misinformation) from the object” ([1982](#), 121, footnote omitted). Evans calls these *information-based thoughts* and provides a specific account of how their referents are fixed.

He begins with the notion of an information system, which is a mechanism that operates on an input to produce end states which somehow encode content. This content takes the form of open predicate formulas, e.g. $F(x)$ or $\text{Red}(x)$. Evans gives photographic cameras as his prime example of such a system. He is silent on how the content of the output states need be encoded; presumably the reader can fill in their preferred theory (e.g., [Dretske 1988](#)). What's important for Evans is that the information output by these systems is about the object input into the mechanism. For instance, a picture is a picture of the thing in front of the camera when it was shot ([Evans 1982](#), 122–29). Evans is clear that—when the mechanism is working—the output encodings should reliably covary with the actual properties of the input. Output accuracy is determined by the extent to which it matches this input.

If a thought is information-based, the thinker's dispositions to “appreciate and evaluate” the thought's truth are “a causal consequence” of their acquisition of information through such a system ([Evans 1982](#), 122). Evans reasonably notes ([1982](#), 138) that in this case “the overriding point or purpose of the subject's thinking is to be thinking of the object from which the information derives.” Hence, it makes sense to talk about the thought as “aiming” at the input of the information system, so that ([1982](#), 139) “Finding the *target* of an information-based particular-thought would involve tracing back the causal routes by which the relevant information is derived from the relevant object.”

But for Evans the target of an information-based thought isn't necessarily its referent. According to Evans, thoughts are mediated through "Ideas" (see 1982, 104). It's unclear if we're to think of Ideas as abstract thought-content constituents (e.g., Fregean senses), mental state types (e.g., singular terms in a language of thought), or physical representation vehicles in the head. Evans compares Ideas to concepts of properties and notes that they allow us to "think of an object in a series of indefinitely many thoughts, ... in the same way" (1982, 104). Also important is that these Ideas involve modes of identification. This is important because Evans accepts what he calls Russell's Principle, the idea that to think of an object "one must *know which* object it is that one is thinking about" (1982, 65, footnote omitted). For Evans, the target of an information-based thought is its referent only if it's also picked out by its mode of identification (1982, 139). In general, an object is a referent of an information-based thought iff it's both the target and picked out by the mode of identification.

Evans conceives of demonstrative thoughts as characterized by a demonstrative mode of identification (Evans 1982, 143–45). He says that demonstrative thoughts are information-based thoughts in which the relevant information system is a sensory system. In demonstrative modes of identification the identifying information is spatial: knowledge of where the object is located. We identify *which* object we're thinking of via spatial information delivered from the sensory information link. This information, he says, needs to fully distinguish the object from all others, and so must be information about its location framed in "public" spatial coordinates (1982, 162); by this he seems to just mean non-egocentric coordinates. What perception delivers, he says, is purely egocentric spatial information. So demonstrative modes of identification identify objects by specifying their locations in egocentric space, but this suffices because of our general background knowledge of how to transform egocentric space into public space (Evans 1982, 165–71).

One might worry that demonstrative thoughts aren't *really* relational, on Evans' account. Isn't the referent fixed by the mode of identification within the demonstrative Idea, which specifies the referent's location, instead of by the information link itself? Evans' response is not clear, but the idea seems to be that information-based modes of identification depend on the information link, in some important way, to successfully fix a target or to have any content at all (1982, 139–41). For example, if my mode of identification doesn't pick out the target of my thought, I can't still be ascribed a thought with the content encoded in that mode of identification, since my thought (if any)

was aiming at the target. So modes of identification don't fix referents; they're merely necessary for knowing which thing is being thought about. The referent, if any, is the target at the end of the information link.

4 A Counterexample: Stimulus-Incorporating Dream Experiences

The problem with Evans' account—the first of two—is that there are cases in which a thinker's demonstrative mode of identification fails to specify the spatial location of a thought's target, but in which the subject still seems, intuitively, to successfully think the demonstrative thought. These are cases of demonstrative thought based on dream experiences incorporating real sensory stimuli. Stimulus-incorporating dream experiences are generally overlooked in philosophical theorizing. Jennifer Windt (2018) puts them to work developing a new view of the nature of dreaming. While my discussion draws heavily from Windt's work,⁶ Windt's more speculative and controversial claims won't be needed for most of my argument.

The standard view has it that dreams are an envatted state in which the brain is cut off from the outside world through sensory and motor blockades of incoming and outgoing neural signals (Hobson 2009, 809; Windt 2018, 2587). On this view dream experiences result from activity within sensory cortex that propagates free from external influences. For example, recent work within the predictive processing framework (e.g., Clark 2013) has it that, while dreaming, incoming sensory signals are lost and the brain's predictive modeling proceeds without constraint, resulting in the typically unstable, often bizarre experiences we enjoy in dreams (see Clark 2012; Hobson et al. 2014; Hobson and Friston 2014).⁷ If this view is right, then what's experienced in a dream is an hallucination in the classic sense that it's an experience of what's not there. According to Evans (1982, 173, 199), demonstrative thought based on hallucination is impossible, since there's no object targeted by the information link. Perhaps more accurately, there is no in-

⁶Windt's work on these cases brought them to my attention. Her work on stimulus-incorporating dreams serves as an excellent touchpoint because of the comprehensive and philosophically insightful way in which she engages with the empirical literature.

⁷The core idea of predictive processing accounts is that experience-generating sensory activity results as the brain tries to construct simulations of the outside world which predict incoming sensory signals.

formation link at all. So, the standard envattment view entails that dream experiences do not allow for (demonstrative) selection of what’s experienced, for the simple reason that (if correct) there is nothing there being experienced to be selected.

While it’s undeniable that much of our dream experience is hallucinatory and that dreams are substantially “envatted” via sensory and motor blockades, the idea that *all* dream experience is cut off from sensory stimuli is gainsaid by familiar—if philosophically underappreciated—phenomena. For example, some people report that occasionally the sound of their alarm will be incorporated into a dream before waking them up. This might take the form of the alarm’s actual sound becoming integrated into the dream, so that you hear it itself among the rest of the dream experience, or the sound may be distorted and integrated into the dream as a different repeating object of (presumably auditory) experience.⁸

Windt catalogues a number of other integration cases studied by psychologists. She notes (2018, 2588) that “there is evidence that different types of stimuli (such as light flashes, sounds, or odors) can be incorporated in dreams, but the overall incorporation rate tends to be variable and quite low.” She summarizes:

The highest incorporation rates appear to be achieved for body stimulation. For example, in one study, sprays of water on the skin were incorporated in dreams in 42% of cases; by contrast, sounds and light flashes were incorporated in only 9 and 23% of cases, respectively (Dement and Wolpert 1958). Inflation of a blood pressure cuff on the leg leads to particularly high incorporation rates between 40 and 80% (Nielsen 1993; Nielsen et al. 1995; Sauvageau et al. 1998). Most instances of incorporation were directly related to the blood pressure cuff or more generally to the legs, as in dreams of wearing strange shoes, having trouble walking, or even experiencing pain. Typically, blood pressure cuff-related themes also seemed to be a motivating factor driving the dream narrative. In other reports, the stimulus was displaced to the non-stimulated leg or even to other dream characters. (Windt 2018, 2588)

⁸Windt (2018, 2587) mentions this case, along with the similar case of a siren being integrated, although I believe it’s fair to say this is a fairly well-known folk phenomena.

Following up on the rich way bodily sensations can be incorporated into dreams, Windt (2018, 2588) notes that there “is also some evidence that vestibular stimulation during sleep using rotating chairs (Hoff 1929; Hoff and Pötzl 1937) or rocking (Leslie and Ogilvie 1996) leads to vestibular sensations such as flying in dreams”. She further speculates (2018, 2588) that “The feeling of extreme effort, incomplete control over the legs, or outright paralysis [in dreams] can plausibly be explained as an illusory perception of REM-sleep-related muscular paralysis—almost as if the relative immobility of the sleeping body were hampering internally experienced dream movement.” A last example she gives (2018, 2588–89) is that “At sleep onset, when the feeling of the bedsheets on one’s skin or the weight of one’s body on the mattress are lost, this can result in vestibular sensations of weightlessness, floating, or falling (Cheyne 2003).”

Starting with these preliminary examples of external-stimuli incorporation, Windt notes the wide range of body-based stimulus incorporation and asks whether all bodily experiences (e.g., tactile experiences, visual experiences of the body, and movement sensations) in dreams are in fact generated by uptake from actual body stimulation. Speculatively, Windt suggests an affirmative answer and sketches the following view. She begins with the work of Mark Blumberg (Blumberg 2010, 2015; Blumberg et al. 2013) on how the sleep motor blockade lets out select motor commands and proprioceptive signals caused by the resulting muscle twitches come through the sensory blockade (Windt 2018, 2590). This process is what Windt calls *bodily self-sampling* and (as Blumberg proposes) allows the brain to construct more accurate body and action maps by attempting to predict the effects of the outgoing commands and adjusting these maps to eliminate the discrepancy with the proprioceptive feedback. According to Windt’s speculative suggestion, bodily experience during dreams—including tactile sensations and proprioceptive awareness of limb position—results from these attempts to model the feedback of this bodily self-sampling, along with the sorts of other purely incidental external stimulation mentioned above (Windt 2018, 2616–18). Of course, this fairly limited combination of internal and external stimulation couldn’t account for the sorts of rich bodily experiences we enjoy while awake, but Windt also argues that dream bodily experiences are much more sparse and indeterminate than their waking counterparts (2018, 2583–87). Overall, Windt proposes a body-centric view of dreams on which dreams (or at least a subset of them) are driven by bodily sensory input around which the brain attempts to model the actual body, filling in an external world in other sen-

sory modalities (e.g., vision), using cognitive resources like memories, which might fit that body model.

While Windt’s full view will provide material for further speculation later in this paper (section 8), for now what’s important is just the minimal claim that there are clear cases of stimulus incorporation in dreams. What’s important about these cases is that the resulting dream experiences are plausibly *not* hallucinatory. Windt is focused on the case of bodily experiences and claims that dream bodily experiences are illusory experiences of your actual sleeping body (Windt 2018, 2590). For example, if you experience yourself running in a dream, you might not be hallucinating a nonexistent running dream body, but instead might be suffering an illusory experience of your own sleeping body running (one based on actual proprioceptive feedback of twitching leg muscles). A similar case can be made for other sensory modalities as well. For example, if actual auditory stimulation from the sound of your alarm or a passing siren causes you to hear (as would be natural to say) that sound from within your dream, there seems to be no reason to think that your experience of the sound is an hallucination. If you experience the sound with little alteration, there doesn’t even seem to be reason to describe the experience as defective or illusory—it simply turns out to be a case of successful veridical perception from within your dream. Describing the experience as hallucinatory merely because it’s incorporated into a dream, despite its faithfulness and causal history, would require a convoluted interpretation of the facts.

It’s worth noting that actual stimuli and dream experiences can sometimes be locked in especially tight causal couplings. Consider the following dream self-report from Tore Nielsen (Nielsen 1992, 360) (specifically, a report of hypnagogic imagery collected while seated), cited by Windt (2018, 2612, see also 2594–96):

Someone in front of me is doubled over towards me, praying.
Someone else reaches around from behind this person and quickly lifts him into an upright position. At the same time I feel my head nodding forward and it awakens me.

In this report the dream involves real-body feedback (the proprioceptive input from a nodding head) affecting the visual component of the dream (transferred onto a non-self dream character), along with dream imagery correlating with (or perhaps affecting) actual body position (as the head goes back up). The coupling is so tight that Nielsen’s dream body experiences the head nod

at the same time as observing the other figure bow forward. Given this tight causal coupling, the simplest interpretation is that while many aspects of dream experiences are straightforwardly hallucinatory, other aspects really are experiences of incorporated internal and external stimuli. Hence these incorporations are not hallucinations, but instead are experiences which make genuine perceptual contact with the world (albeit ones which often heavily distort or alter the stimuli with which they put one in contact).

Now we can return to Evans and the question of whether dream experiences enable the demonstrative selection of what's experienced for thought. As noted above, when the dream experience is hallucinatory the answer is negative. But there will be many dream experiences (or aspects of the overall dream experience) which are not hallucinatory and put us in genuine perceptual contact with the world. While dreaming, your own sensory systems (especially your proprioceptive and other interoceptive systems) still function, albeit in an attenuated or nonstandard way. For example, your body-monitoring systems function to facilitate bodily self-sampling, not to facilitate sensorimotor action with the world, and your exteroceptive sensory systems are largely (but not entirely) suppressed. Still, these systems provide an information link—or at least a perceptual link—to your body and (to a lesser extent) the outside world. Through this link your body and external stimuli can become targets of demonstrative thought. So, you could attempt to think various demonstrative thoughts, based on the dream experiences which manage to make genuine contact with your body or the external world.

For example, an intruding alarm or siren could be the target of an audition-based demonstrative thought.⁹ With respect to the body, the most natural examples would be what Evans calls self-ascriptions (1982, chapter 7), e.g. thoughts like “*my legs* are moving fast”, “*I* am standing”, or “*my right foot* is in such-and-such position”. For Evans these self-ascriptions are a variety of demonstrative thought (1982, 205, 212), and so are germane to the discussion. But you could also attempt bare demonstrative thoughts based on interoceptive bodily experience without the self-referential conceptualizations making for self-ascriptions. For example, while attending to your proprioceptive experience of your right foot in a dream, you could attempt to think “*that* is an awkward position”.

⁹Sounds, I take it, are potential objects of demonstrative thought (see [O’Callaghan 2007](#); [O’Callaghan 2016](#)).

According to Evans (1982, 253), for these attempts at self-ascriptions and other demonstrative thoughts to succeed, it's not enough that you have a functioning information link to your body or the other incorporated external stimuli. You also have to be able to identify the spatial location of the relevant stimulus, be it a sound, body part, body position, or something else (although you need not in fact identify the location). But I claim that (a) the balance of intuitions in these cases favors saying that these demonstrative thoughts in the dream succeed, and (b) in these cases you are not in a position to identify the locations of the thoughts' targets in public space. If (a) and (b), then Evans is wrong that demonstrative selection requires knowledge of—or the ability to find out—the target's location in public space.

To motivate claim (a), imagine such a case. You are walking along in a dream, and within the dream you have a proprioceptive experience as of the position of your right foot. It is, suppose, a vivid experience which would be hard to introspectively distinguish from a normal waking proprioceptive experience of your foot in that position. You attend carefully to what you experience as your foot and its position and attempt to think “*that* is an awkward position”. Assuming your experience is based on an actual *perceptual* link to your foot (it is merely an illusory experience, not an hallucination), your attention is enough to succeed in thinking this demonstrative thought. Alternatively, imagine a case of especially tight auditory integration. As you dream, whatever else may be happening around you, you begin to hear the characteristic sinusoidal pitch changes of a siren. Although you can't be sure of the source from within the dream (and likely don't even ask yourself about it), the sound actually is from a passing ambulance. You attend to the sound and think “*that* is a siren”. Again, given that you actually are hearing a real sound, you succeed in selecting the sound and thinking the thought.

Claim (b) is straightforward. That you are unable to locate your thought's target within public space follows from the nature of dreaming. Although the specific experiences in the above examples (of your foot and of the siren) are perceptual experiences which make contact with the world, most of your dream experience is still an hallucinatory stimulation of a nonexistent dream world. While you have a few minimal surviving perceptual links to the outside world (e.g., to your foot or to the siren), the kind of robust spatial perceptual contact you have with the world while awake which tells you about the public space around you is gone. Instead of your brain functioning to reconstruct the space around you, as Windt suggests it seems that your dreaming brain uses your exteroceptive sensory machinery (especially vision)

to simulate an imaginary world in which your dream body and actions are situated. In effect, while dreaming you are cut off from how your egocentric space maps onto public space, as your brain generates a map of public space not based on actual sensory input, but based on a simulation fitting your proprioceptive feedback.¹⁰ Just as a fictional literal envatted brain lacks the kind of informational access to the public world needed for knowledge of public location (Evans 1982, 250–53), you during a dream lack the right sort of access too. While dreaming you lack the ability to know the location and position of your foot in public space, or to know the location of the siren, because the world (but not your foot or the sound) has been replaced with a radical hallucinatory simulation.¹¹

What about Russell’s Principle and the idea that successful thought requires knowing which object one is thinking about? The phrase ‘know which’ is slippery. I’m suggesting that we shouldn’t understand it in the demanding sense proposed by Evans; the knowledge-which encoded in demonstrative modes of identification isn’t knowledge of public spatial location. But there is a sense in which we still know which thing we’re thinking about in the dream examples. We know which item we’re thinking about—the sound, our foot, or foot’s position—via our perceptual experiences of that item. Which item is it? It’s the perceptually experienced one I’m attending to now as I think the thought. Given that my experience puts me into contact with the thing, this amounts to a kind of knowledge of which item is at issue.

Evans gives examples in which a subject has an information link, but demonstrative thought is impossible (e.g., 1982, 149, 165) and the problem (he suggests) lies in an inability to locate the object in public space. He

¹⁰Even if Windt’s positive account of what the brain is doing is not correct, what is certainly true is that while dreaming your brain is no longer in the kind of sensory contact with the environment needed to generate a map of public space.

¹¹Tyler Burge (2010, 200) and Christopher Peacocke (1983, 170) also have given purported counterexamples to Evans’ spatial localization requirement. These examples are things like seeing an object in a mirror without noticing the mirror and seeing a star the light of which is refracted in the atmosphere. The problem with these examples is that Evans only requires that the information link *enable* locating the thing in public space. In these sorts of cases I may not in fact locate the object correctly at first glance or based only on a momentary slice of my experience, but perception still affords the kind of information link which would enable me to locate the object after some exploration. But in the dream case my information link to my body won’t allow me to locate my body parts *in public space*, in principle, because while I’m informationally connected to my body, I’m not informationally connected to public space (to the distal environment).

also has an argument against the sufficiency of an information link (1982, 148–50). I agree with Evans that a sensory information link is not sufficient for demonstrative thought. But the problem isn't that there is some other condition required as well—e.g., knowledge of location in public space. The problem is that information links were never really what was important for demonstrative thought in the first place. What matters is a perceptual link (successful perception of the target), and not all sensory information links afford genuine perception. This is the second problem with Evans' account: Evans misunderstands what it is about perceptual links which makes them exploitable for demonstrative selection. Evans' emphasis on information links and his posit that they are fundamental referential links is wrong.

5 Information Links

Evans' idea that information links play a central role in relational thought has been highly influential. Much of what is to follow will be an attack on this idea. But it's worth first articulating a modern version of the view. François Recanati says:

In perception, we are related to the object we perceive. The perceptual relation is what enables us to gain (perceptual) information from the object. ... In general there is acquaintance with an object whenever we are so related to that object that we can gain information from it, on the basis of that relation. Acquaintance relations are *epistemically rewarding (ER) relations*, on this view. ... To think of an object directly or non-descriptively is to think of it through *some such relation*. In such a case, what determines the reference—what one's thought is about—is the relation: the reference is the object to which we stand in the relevant relation, ... (Recanati 2012, 20, footnote omitted, emphasis added)

In *Mental Files* (2012), Recanati develops his own version of this view, based on a mental-file approach to nondescriptive, i.e. singular, thought.¹² Mental

¹²Evans' work is an example of the mental file approach as well, although he doesn't use the term 'mental file'. As noted in footnote 4, singular thoughts are any thought "whose content is a singular proposition—a proposition involving individual objects as well as properties" (Recanati 2012, 5, footnote omitted). Although Recanati seems in this quote to assume that all singular thoughts are relational, some have argued otherwise (e.g.

files themselves have a long history. The basic idea is that, within the mind, collected information is sorted into a file-like structure. A mental file is some subset of stored information that's treated as if it's about the same thing (see [Dickie 2010](#), 222).

According to Recanati (2012, 34–5), “The role of the files is to store information about the objects we bear these [epistemically rewarding] acquaintance relations to.” These files are “the mental counterparts of singular terms” (2012, 35, footnote omitted). He says (2012, 59–60) that files “are typed according to the type of ER relation they exploit”, the associated type of ER (epistemically rewarding) relation being determined by the functional role of the type. As he says (2012, 60–61, footnote omitted), “Since the function of a (type of) file is to exploit a given (type of ER) relation, a token of that type should come into existence only if the subject stands in the appropriate contextual relation to some entity, a relation in virtue of which it will be possible for him or her to gain information from it.” Applying this account to demonstrative thought, he says (2012, 62) “demonstrative files, such as the files THAT MAN or THAT THING, are based on certain contextual relations to objects, in virtue of which we can not only perceive them but also focus our attention on them in a discriminating manner”.

To summarize, Recanati's view is that we think relational thoughts by tokening mental files. The referent of these thoughts is the referent of the mental file (see also [Dickie 2010](#), 222; [Jeshion 2010](#), 129), which is the thing with which we are acquainted through the epistemically rewarding relation (the information link) feeding the file—the target, as Evans would have said. For demonstrative thoughts, the relevant mental file is of a type which collects information based on attention to a consciously perceived item. If this account is correct, then it's a structural feature of demonstrative thoughts (and relational thoughts more broadly) that information links play an essential role in fixing referents: those thoughts are made possible by tokening mental files, which by their nature track and refer to things through information links.

As noted above, this account makes perception *incidental* to demonstrative selection. What's really important in demonstrative thought—being a type of file-based singular thought—is that there is an information link feeding a mental file (see also [Smithies 2019](#), 66 for this interpretation). Attention

[Jeshion 2010](#)). To avoid this complication, I will discuss Recanati's account only with an eye towards relational thought and not his broader topic of singular thought.

to a consciously perceived object just so happens to be what constitutes that link in this case, but what does the reference-fixing is the epistemic nature of that attention: that it provides a feed of information which is filed together. Recanati (2012, 35, emphasis his) emphasizes this priority: “The characteristic feature of the relations on which mental files are based, and which determine their reference, is that they are *epistemically rewarding* (...).” He continues (2012, 35): “Relations of perceptual acquaintance are ER relations: they are the sort of relation to objects which makes the perceptual flow of information possible.”

6 Are Perceptual Links Necessarily Information Links?

Like Evans, Recanati’s account is elegant and compelling. Both put information links at the center of demonstrative thought: we select consciously perceived items for thought by exploiting the information links opened up by our attention to those items. That these links are also perceptual is not fundamental on these accounts; they entail that if the links were not also information links, the links would not have enabled demonstrative selection. What I suggested at the end of section 4 was that this approach is wrong: perceptual links need not be information links to enable demonstrative selection. That typical perceptual links happen to be information links is merely incidental.

Before getting to my main argument against the information-link approach (in sections 7 and 8), I want to consider a possible first reaction from its proponents. This reaction goes that perceptual links are *necessarily* information links, and so there are no cases of perceptual links that are not information links. Exploring this issue will clarify what’s at stake for the information link approach.

Why think—as this reply proposes—that all perceptual links are information links? You might suggest that if a sensory experience, or the associated underlying causal connections between distal stimulus and neuro-psychological state, failed to constitute an information link, then they would also fail to be a genuine perceptual link. After all, many hold that perceiving an object requires discriminating it from the background (Siegel 2006; Schellenberg 2019), tracking it through changes in proximal stimulation (Burge

2010), or locating it in egocentric space as a target of action (Matthen 2005, chapter 13)—all information-hungry activities. In addition, many philosophers think that there is something inherently content-bearing about perceptual phenomenology, so that, e.g., there’s an information-conveying way things necessarily appear in experience (e.g., McDowell 1994; Siewert 1998; Siegel 2010; Johnston 2011).

But this move is too fast. It does not follow from perception being somehow tied up with the flow of information (as it surely is) that perceptual links necessarily constitute information links *in the sense envisioned by Evans and Recanati*. There are two features characteristic of information links, as described in their accounts: (1) they are accurate when the underlying mechanism functions normally, and (2) they are exploited to select targets of thought by tokening the mental files they feed. These information links are *epistemically robust* in the sense that they provide accurate information when things are going well and provide usable information to the subject by feeding into thought-constituents (mental files) deployed by subjects (and are exploited by this deployment).

Evans makes (1), the requirement for accuracy, explicit (see Evans 1982, 125). Recanati is less clear on this accuracy requirement, but his emphasis on links being *epistemically rewarding* suggests it as well. Other information link proponents, including Dickie, endorse the requirement. Regarding feature (2), as Evans and Recanati conceive them, information links provide information that’s usable by the subject. For Evans, these links funnel information into a subject’s “controlling conception” of the target in a way that causally affects their beliefs and judgments (Evans 1982, 122). For Recanati, these links funnel information into a mental file. Recanati explicitly takes mental files to be constituents of the thoughts of rational subjects; they are “non-descriptive mode[s] of presentation” of the kind figuring in explaining the informativeness of identity statements (Recanati 2012, 34). This view of mental files as they figure in relational (and singular) thought is fairly common (e.g., see Jeshion 2010, 131).

It should be fairly clear that perception itself could be thoroughly wrapped up in the flow of information—as posited by the various theories of perception cited above—without meeting either of these requirements. Take, for instance, Mohan Matthen’s view (2005, chapter 13) that perceiving a particular object requires that visual processing locate that object in egocentric space as a target of action. There is no reason to think this locating-for-action further requires visual processing to feed into the mental files stocking our

general conceptual store, that it provides a link which we exploit by tokening a mental file from this conceptual store, or that it requires any accuracy within any of our visual states. Similarly, tracking distal objects and properties through changes in proximal stimulation (Burge 2010) is something our perceptual systems can do at an entirely subpersonal level or without much accuracy. Likewise, other mainstream theories of perception do not entail that perceptual links are *epistemically robust* information links.

I will argue in the rest of this paper that there are cases of perceptual links which fail one (or both) of these requirements, and so fail to be epistemically robust information links, while still enabling demonstrative selection. Stimulus-incorporating dreams, as I'll discuss in section 8, fail the accuracy requirement. I'll argue in the next section (section 7) that cases of multimodal sensory binding failures involve perceptual links which cannot be exploited by tokening the mental files they feed.

7 The Multimodal Binding Case

Imagine that you see a ball rolling across a table and hear a sound you mistake as the sound of the ball rolling. The mistake isn't just at the level of deliberate judgment, but is a perceptual binding error. You mistakenly hear the sound as coming from the ball, in the same way you mistakenly hear a ventriloquist's voice as coming from their doll or mistakenly hear, while watching a movie, the sound of a character's voice as coming from the image on the screen.¹³ What you hear is in fact the sound of a steel bearing rolling across a table behind you. Imagine you voluntarily (i.e., deliberately) focus attention on just the noise, while the ball stays in view. It is intuitive that in this case you could, with referential success, demonstratively think of the attended sound in a thought like "*that's* loud".

But this is not the result predicted by Recanati's account. Since the mistake is a *binding* error at the level of perception, there's a single perceptual file collecting information from both the seen ball and the heard noise.¹⁴

¹³Often (if not normally) audio-visual binding is veridical, as when you hear the speech of someone you see in person as coming from them, or hear a screech as coming from a car you see skidding to a stop.

¹⁴I am not presupposing any account of the files within our sensory systems, but perhaps the best known work on perceptual files is from Daniel Kahneman and Anne Treisman (Kahneman and Treisman 1984; Kahneman et al. 1992; Treisman 1998).

This is because when your sensory systems bind information together they are collecting information into a single perceptual file. The perception-based mental file involved in this case—the one you would token to think thoughts about the target at the end of the perceptual links feeding it—is not just that perceptual file from your sensory system, but it includes all the information collected in the perceptual file, unless you explicitly recognize the binding as an error. Just as the perceptual file is fed by both auditory and visual links, so is the mental file you deploy to think the demonstrative thought. So, Recanati is forced to say that in this case the referent of your thought (if any) is the referent of a mental file drawing on both your visual and auditory perceptual links. The problem is that this forces Recanati to say one of a few things, none of which gets the sound as the referent.

First, he might say that the lack of a unique information source leads to reference failure. This is a common thing for information-link theorists to say about cases in which (unknown to the subject) the item they are tracking is switched (see [Evans 1982](#), 170–73; [Recanati 2012](#), 132). In this case there is no diachronic item swap, but the multiple synchronic information links to distinct things might yield the same result. But Recanati shouldn't say that the relevant mental file fails to refer, since you are able to think demonstratively about the attended sound.

Second, he might say that the referent is the dominant causal source of information. This is what Recanati says about what he calls “encyclopedic files” or “encyclopedia entries” ([2012](#), 46, 73, 140), which are files that collect information from various channels over time on (what the subject takes to be) a single object. The problem is that on a flat-footed reading of this idea, it is the rolling ball, not the sound, which is the dominant causal source of information.

The standard measure for quantifying information has it that information quantity has to do with eliminating alternative possibilities ([Shannon 1948](#)). Your visual experience of the rolling ball eliminates a lot of alternative possible situations: situations in which the ball is some other color, has a different shape, a different position relative to the table, a different speed, a different trajectory, and so on. Your auditory experience of the sound of the rolling steel bearing will likewise rule out certain possibilities—the possibility that the bearing is a shape that won't roll, the possibility that the bearing is stationary, etc—but the visual experience rules out all these same possibilities about the ball plus many more. Hence, more information is collected about the seen ball, which is thus the dominant causal source.

Recanati doesn't have in mind this flat-footed construal of 'dominant causal source'. Instead he cashes out the notion in terms of "activated regions" of mental files (Recanati 2012, 138–41). The idea is that when files are fed by multiple information links (e.g., as with encyclopedic files), stored information from each link forms a distinct "region" in the file and these regions can be "activated" to different extents. Recanati suggests that the referent of a file (at a time) is the thing linked to by whichever region is most activated. He also suggests that the dominant causal source of information is determined by whichever file region is most activated.

Recanati might press that this idea of differentially activated file regions can be applied beyond encyclopedic files to perception-based files to solve the problem.¹⁵ Because you're attending to the sound, the region of your mental file populated by auditory information is more activated than the visual region, and hence the referent of your thought is the sound.¹⁶ I worry that this solution fails without a more robust explanation of what's meant by 'file regions' and 'region activation'—notions which Recanati doesn't explain and which stress the already metaphorical talk of 'files'—but there is a more pressing problem.

First, the case is complicated because we have file systems at multiple levels of mental architecture: we have the files within our sensory systems along with the mental files which constitute post-perceptual thoughts (such as our demonstrative thoughts). What's relevant for this potential solution are regions within a (post-perceptual but still perception-based) mental file: if your demonstrative thought about the misbound sound is explained by tokening a (post-perceptual) mental file, then there must be a region within that file, fed by just the auditory information link to the sound, which can be activated. The pressing problem is that we have no reason to think this mental file has distinct regions for the various modalities feeding it.

¹⁵This would be a departure from Recanati's theory, which treats encyclopedic files as akin to proper names (Recanati 2012, 46). Unlike perception-based files which are supposed to be tied to a single perceptual link, encyclopedic files are maintained even as any one link feeding the file is lost. New links are hooked up to one of these files as the targets of those links are identified as being the same as the target of the file. Recanati hasn't considered the possibility that a perception-based mental file might itself be fed by multiple perceptual links which might diverge in their target.

¹⁶This response is based on the one Recanati gives in (2012, chapter 11) to how one could successfully think singular thoughts (after they've been unknowingly transported to Twin-Earth) about Twin-Earth doppelgangers of people they know, given that they're thinking of them through a mental file which draws on information links to different people.

Because Recanati himself doesn't consider the dynamics involved in multiple (synchronic, multimodal) perceptual links we left to speculate,¹⁷ but there are two reasonable ways in which a perception-based mental file (i.e., a file fed by perceptual input which can be deployed at the level of thought) could come to have distinct regions, corresponding to distinct perceptual links, within it: (1) if it was being fed by multiple distinct perceptual files each tracking one of those links, or (2) it was fed by one perceptual file (itself fed by multiple perceptual links), but the subject recognized the binding error and so was able, at the cognitive level, to distinguish different sources of information. By stipulation, in the case under discussion you don't realize the sound you hear isn't coming from the ball you see, so (2) is out. But as discussed at the start of this section, binding errors are cases in which sensory systems slot incoming information into the same perceptual file (i.e., treat the information as about the same thing), so the perceptual file directly available to cognition is one fed by both the auditory and visual links. So option (1) is out as well.

Is it really the case that your sensory systems only produce (and make available to cognition) a single perceptual file on the misbound ball-sound pair?¹⁸ As Casey O'Callaghan discusses in his work on multimodal percep-

¹⁷Recanati tends to talk (incorrectly) about perception as if it provides a single, clean information link, or as if perceptual information links will map one-to-one to the perception-based mental files supporting demonstrative thought (e.g. Recanati 2012, 35, 62).

¹⁸It might further be suggested that distinct auditory and visual files are formed which your sensory systems merely "link" together; the linking would explain the apparent perceptual binding and the two distinct perceptual files would afford differentially activated regions within a post-perceptual mental file. File linking is an operation that Recanati posits at the level of mental files (2012, 42–47) as a way to explain the cognitive significance of identity statements and beliefs. Recall that for Recanati mental files are something like, or correspond to, modes of presentation. If you believe an identity holds between the referents of two such modes of presentation then there must be some cognitive operation which preserves the distinct modes of presentation while explaining how and why you easily treat information in each file as pertaining to the referent of the other file. Linking (understood as less extreme than merging files) is the operation Recanati posits to explain how modes of presentation are preserved while also allowing information transfer. As suggested by an anonymous reviewer, Recanati might try to extend the linking operation as a way to handle cases of multimodal binding failures. The first worry for this approach is that the sound and the seen rolling ball are *not* identical and not experienced as identical, so it wouldn't make sense for the perceptual system to keep a file on each and link the two, if this linking corresponds to representing identity. It might be suggested that what's linked isn't the sound file and the file for *the ball*, but instead the sound file and a file for the seen *event* of the ball's rolling. The ball's rolling might be an event perceived through

tion (2019), empirical evidence demonstrates deep crossmodal interactions in sensory processing.¹⁹ For example, in the misbinding case under consideration here, hearing the sound as coming from the seen ball involves, in part, vision modulating the spatial content of audition: although the sound is coming from behind you and the proximal auditory stimulation encodes that location, sensory processing recalibrates that location on the basis of visual information so that you hear the sound as in front of you (coming from the seen ball). There are many other crossmodal interactions which come out through illusions (see O’Callaghan 2019, chapter 2.1). Two well-known cases are the sound-induced flash illusion in which a single flash accompanied by a double beep is seen as two flashes, and the McGurk effect in which the /ba/ sound, accompanied by video of a speaker saying /ga/, is heard as /da/ (ibid). As O’Callaghan discusses (2019, chapter 2.2), there is ample evidence that these effects happen within perceptual processing (and so are not post-perceptual cognitive modulations): subjects don’t report any reasoning on their part, the effects happen fast and automatically even in naïve subjects, and disruptions to sensory cortex through methods like TMS (transcranial magnetic stimulation) modulate these effects. O’Callaghan further notes (ibid) that, specially in the sound-induced flash illusion, fMRI (functional magnetic resonance imaging) shows that the auditory stimulation actually modulates primary visual cortex. Finally, as O’Callaghan notes (ibid), there are sensory neurons which are superadditive with respect to input from different sensory modalities. These superadditive responses encode the joint presence of multimodal stimuli and show neural sensory processing of input from multiple modalities.

What this evidence shows is that even if our sensory systems did produce distinct modality-specific representations of distal stimuli (e.g., modality-specific perceptual files), those modality-specific representations would not draw on modality-specific perceptual links. Hence, even if there were distinct

both vision and audition. While some version of this response will likely make the linking reply workable, there are two more basic problems with this overall linking approach. I’m about to argue that even if there were distinct visual and auditory files, each of those files is fed by links which make contact with both the sound and seen ball. More importantly, as I’m also about to argue, there is good evidence that your sensory systems keep a *single* file in these cases, instead of producing separate files which could be linked.

¹⁹O’Callaghan’s work on multimodal perception (especially 2015; 2019) is what drew my attention to these cases and the complications they cause for theorizing about demonstrative thought.

auditory and visual files on (respectively) the rolling sound and the ball, each of these files would be drawing on both auditory and visual information. So it seems unlikely that, even at the most basic levels of sensory processing, we will find modality-specific perceptual files with a perceptual link to only the sound or the ball.

But as O’Callaghan discusses, there is reason to think that, in cases of multimodal binding, there are not distinct modality-specific perceptual files. First, feature binding in the intramodal case (e.g., binding seen color and seen shape) is generally taken as evidence of a single perceptual file, so it’s unclear why this shouldn’t be extended to the intermodal case (e.g., binding sounds with seen shapes) (O’Callaghan 2019, chapter 3.2.2). Second, object-specific preview effects have been observed in the intermodal case. As O’Callaghan explains:

Zmigrod et al. (2009, 674–5) report that patterns of performance that characterize unimodal feature binding occur intermodally between audition and vision, and between audition and touch. For instance, based on a study presenting a sequence of colored circles, each paired with a specific pitch, they argue that color–pitch pairs may be bound (each stimulus was presented for 50 milliseconds, the second separated by 450 milliseconds from response to the first). That is because presenting $color_1$ with $pitch_1$ at t_1 followed by $color_1$ with $pitch_2$ at t_2 impairs responsiveness at t_2 (slows it and degrades its accuracy) both to $color_1$, when compared with repeating $pitch_1$ in the color identification task, and to $pitch_2$, when compared with pairing distinct $color_2$ in the pitch task (Experiment 1). The fact that an auditory pairing incurs object-specific preview effects for vision conflicts with what is predicted by the hypothesis that we harbor only modality-specific object files with intramodal binding. (O’Callaghan 2019, 63–4)

The idea is that if feature binding across modalities wasn’t indicative of a single perceptual file collecting both the auditory and visual features, then these preview effects would not be observed. Finally, the spreading of (involuntary) attention from one perceived feature to another is a third indication that those features are being collected in a single perceptual file, and O’Callaghan notes that attentional spreading is observed in the multimodal case as well (2019, 64).

To summarize, there is good evidence that in multimodal binding cases, like the ball-sound case under discussion, sensory systems produce a single perceptual file fed by (in this case) distinct auditory and visual links. Any post-perceptual mental file enabling demonstrative thought to either the ball or the sound would need to draw on this single perceptual file, and so would seem to lack distinct regions which could be activated. The fact that modality-specific features (e.g., auditory location) within this perceptual file are actually computed using information from other modalities is a further complication: even if there were distinct modality-specific perceptual files feeding into the post-perceptual mental file which might support distinct subregions, these modality-specific perceptual files themselves would draw on perceptual links to distinct objects (the sound and the ball). The upshot is that, as you see the ball and (mistakenly) hear the sound of the steel bearing come from the ball you see, you have perceptual links to two distinct items. If information links are the kind of thing we exploit for demonstrative selection by tokening mental files fed by those links, then neither of these distinct perceptual links are information links in this sense: the two perceptual streams are too intertwined in our sensory processing and there just aren't distinct files attached to each link (or even distinct file regions).

So while there might be distinct, activatable regions within some mental files—e.g., within files that combine remembered information with current perceptual information recognized as being about the thing remembered (Recanati 2012, 139)—there aren't such regions within mental files based on multimodal perceptual processing. It's been known for awhile that certain sorts of reference-switching cases (like thinking of a Twin-Earth doppelganger of a close friend after unknowingly being moved to Twin-Earth) pose a challenge for mental file views (see Recanati 2012, chapters 10–11). The multimodal binding case is somewhat similar, insofar as it also involves a mental file with links to distinct items. But it keeps things to *perceptual* links, as opposed to the usual involvement of memory in switching cases, thus making itself a problem for mental file accounts of even the limited (seemingly problem-free) case of pure perception-based demonstrative thought. While Recanati's approach to reference-switching cases using differentially activated file regions might work, it doesn't transfer to perceptual cases involving multimodal binding mistakes.

While I've focused on Recanati's specific version of the mental file account, I believe the case developed in this section will present the same problems to any account which holds that we exploit information links for

demonstrative selection by tokening the mental files feeding those links. The basic problem is that there are cases (like the misbound sound) in which we can demonstratively select an item based on our attention to it in experience, but (given the intertwined nature of multimodal sensory processing) there are no mental files fed by a clean information link to just that item. The problem illuminated by these cases isn't one to be solved by a more complex mental-file account (e.g., making use of activated file regions), but a deeper one. Information links—understood as being exploited by the tokening of the mental files they feed—just aren't what's important in demonstrative thought. No such information link exists to the sound of the steel bearing in the case developed here, but yet demonstrative thought about that sound is still possible.

It's worth pointing out an important implication of the above argument. If correct, what the above argument shows is that facts about the targets of voluntary (i.e., deliberate) attention are more basic than facts about the referents of perceptual and mental files. It does seem that we are able to voluntarily attend to the sound of the rolling bearing, but given the integrated multimodal nature of perception we have good reason to think there's no perceptual or mental file for just this sound. Hence, we can attend voluntarily to things even if we have no single file on that thing. While some theorists might favor a reductive theory of attention on which deliberate personal-level acts of attention reduce to the selection of a file by subpersonal attentional mechanisms, it's important to point out that this isn't the only view available. First, attention might reduce to the selection of some other (non-file-based) kind of mental or neural representation by subpersonal attentional mechanisms. Second, personal-level voluntary attention might not reduce to the action of any subpersonal mechanism at all; it might, instead, merely be an organism-level behavior or environmental interaction facilitated by subpersonal attentional mechanisms (see [Allport 2011](#), 25–26). The work of this section shows that one of these two alternative views of attention is correct.²⁰

²⁰As [Allport \(2011, 25\)](#) points out, the term 'attention' itself is ambiguous. There are several things we could mean by it, and nothing I've said rules out the existence of subpersonal attentional mechanisms which select mental or neural representations (including files) for further processing or kinds of attention (e.g., perhaps involuntary attention) which reduce to these mechanisms. All this section shows is that the specific sort of attention involved in attending to the sound of the steel bearing doesn't reduce to selection of files.

8 Dickie and Justification

Imogen Dickie, in *Fixing Reference* (2015a), provides a compelling, theoretically rich recent example of an account which grounds demonstrative selection in file-based information links.²¹ Her account is interesting because it emphasizes the accuracy requirement (see section 6) and develops it into a sophisticated account on which it's the justification-conferring role of information links which does the reference-fixing work in demonstrative selection. Dickie introduces her view, situating it clearly within the information-link tradition, as follows:

Thought about particular ordinary things is engagement in a specific kind of information-marshalling activity. Each of us receives a stream of information—from perception; from utterances made by others—and marshals it into bodies of beliefs we use terms like ‘that’, ‘he’, and ‘Bertrand Russell’ to express. Like any activity, this information marshalling has a goal. Its goal is to secure what I shall call, as a first pass, ‘cognitive focus’ on things outside the mind: we are trying to tune in on objects in the world around us. I am going to argue that the aboutness of our ordinary thoughts *just is* cognitive focus. ... aboutness is what we achieve [when we achieve cognitive focus]. (Dickie 2015a, 1)

Dickie intends this as a general account of thought about particular ordinary objects, not just demonstrative thought. By ‘ordinary’ objects she means things usually made available for thought by perception and testimony, like people and chairs, as opposed to weird stuff like shadows, numbers, and electrons (Dickie 2015a, 23). She frames things in terms of *bodies of beliefs* because this is a mental file account, and these bodies of beliefs are the mental files (2015a, 50, note 17).

For Dickie, ‘cognitive focus’ is a term of art which roughly conveys the idea that one’s belief-forming mechanisms will get things right about a specific object unless something goes wrong. Her official view is framed in this latter way. An opening approximate version goes (Dickie 2015a, 2): “A body of beliefs is about an object iff its means of justification converges on the object, so that, given how the beliefs are justified, the subject will be unlucky

²¹For other examples of work that takes on important parts of the file-based information-link approach, see (Pylyshyn 2007, especially 18–20, 37–39; and Jeshion 2010).

if they do not match the object and not merely lucky if they do.” As an example, she applies this general account to demonstrative thought (2015a, 5): “the belief I express when I say ‘He is asleep’ while looking at my dog is about the dog because, given the way the belief is justified (by uptake from a perceptual link with the dog), I will be unlucky if the dog is not asleep, and not merely lucky if he is.”

Dickie is very clear that on her view what does the reference-fixing work in any thought about a particular is cognitive focus. Regarding relational thought, she says (2015a, 5), “a relation plays an aboutness-fixing role iff it plays a role in securing cognitive focus, which is to say, a role in securing the result that justification for beliefs a subject would express using a singular term converges on some unique thing.” Later, regarding perception specifically, she says:

For example, consider your beliefs ⟨That is round⟩; ⟨That is orange⟩; ⟨That is rolling⟩, formed as you watch an orange rolling along the table in front of you. These beliefs are justified on the basis of your perceptual link with the orange. [My view] says that the beliefs are about the orange *because* it is the unique thing whose properties you will be unlucky to get wrong and not merely lucky to get right when you form beliefs justified in this way. ... *That is why this is the thing the beliefs are about.* (Dickie 2015a, 37, emphasis added)

Dickie’s chapter on demonstrative thought (2015a, chapter 4), after arguing that perception does in fact provide links to ordinary objects, spends the rest of its argumentation showing (1) that a perceptual link “confers justification on perceptual demonstrative beliefs”, and (2) “that this justification converges on the objects our perceptual demonstrative beliefs are about” (2015a, 122).²² I take it as clear that for Dickie, what’s *essential* for demonstrative selection is that attention to a consciously perceived item (a perceptual link) provides a justification-conferring information channel that converges on a unique item; that this link is perceptual is merely *incidental* on her account. As noted, Dickie emphasizes that other information links (e.g., testimony) could just as well serve to realize this convergent justification-conferring mechanism, thereby enabling relational thought.

²²Dickie has a technical definition for ‘converge’, but the details won’t matter here (Dickie 2015a, 52).

I agree with Dickie that perception must actually link us to particular ordinary objects and that these links must converge on these objects if they are to afford demonstrative thought. This much is baked into the idea that attention to consciously perceived items must afford an *exploitable relation* to ordinary objects if we're to understand demonstrative thought as a form of relational thought. What I disagree with is the further suggestion that a perceptual link must confer justification for potential beliefs formed on the basis of it, and that this *justification* must converge on the target in the sense that we are unlucky if our beliefs fail to match the target and not merely lucky if they do.

Dickie is not the only philosopher to endorse a justification requirement for demonstrative thought. Although accuracy or reliability need not amount to justification, the requirement that information links be accurate or reliable (when functioning well), endorsed by Evans, is obviously in this same spirit. Other philosophers have explicitly endorsed a justification requirement. Declan Smithies says (2019, 70; see also 2011, 21) that "You can think perceptual demonstrative thoughts about an object only if you have perceptual information that gives you defeasible, noninferential justification to form beliefs about the object." Although Smithies talks about information links, his commitment to this requirement is less thoroughly enmeshed in the mental-file framework. Presumably a justification requirement on demonstrative selection is something one could endorse while rejecting almost all the substantial claims of a file-based information-link account.

While there's more to say about Dickie's account, enough has been said to get my objection on the table. (I will return to the details of Dickie's account; while Smithies' account itself is rich and worth engaging, a lack of space precludes further discussion.) First, insofar as Dickie's account builds in the idea that demonstrative selection involves exploiting information links by tokening the mental files they feed, the multimodal binding case presented above is a problem for her just as it was for Recanati and Evans. To put it in Dickie's terms, your (multimodal) information link in this case fails to converge on any item, but you still are able to demonstratively select the sound of the rolling steel bearing.

In the multimodal binding case the information link fails to converge, but does provide justification for beliefs (about the attended sound). So this example won't help with analyzing Dickie's (or Smithies') claim about a requirement for justification. But the case of stimulus-incorporating dream experiences does help. Consider again the dream case: you have a dream

which incorporates real bodily sensory feedback into experience of your limbs, or incorporates the actual sound of your alarm. As discussed (section 4), these are genuine perceptual experiences: by attending to your body or the sound as presented in the experience you have a real perceptual link to it. On this basis, it's plausible that you are able to demonstratively select the attended body part and sound. As I pointed out above, you are not in a position to identify the location of the resulting thought's target in public space. But it's also plausible that these perceptual links do not confer any justification. If that's right, then there is no justification requirement in demonstrative thought.

The claim that stimulus-incorporating dream experiences fail to provide justification should be *prima facie* plausible even before we consider any particular account of justification. First, these experiences tend to be illusory and highly distorted. Even when they manage some measure of veridicality (e.g., as with a faithfully incorporated sound), they result from highly attenuated sensory processing and (importantly) a good bit of luck in having made it through the sensory blockade. Of course, illusory experiences still confer justification in normal waking perception. Your visual experience of the well-known Müller-Lyer lines, while viewing them under normal conditions, presumably justifies your belief that the lines are unequal before you learn of the illusion. But stimulus-incorporating dream experiences are systematically illusory and a matter of luck: e.g., I experience my body as upright and locomoting while in fact I'm prone and atonic, or I faithfully hear my alarm only by chance. What I experience as a powerful stride mid-run is really just a minor muscle twitch in my thigh. While stimulus-incorporating experiences in dreams are generated by actual sensory feedback, save for certain cases where by luck a stimulus is incorporated faithfully, these experiences are distorted in ways that make them unreliable. If Windt's account is correct, then it gives us insight into the source of this unreliability. Unlike during waking illusions, during dreams the sensory systems generating the illusory experience aren't functioning in the normal way. During dreams (as Windt proposes) the brain is doing bodily self-sampling, which is functionally very different from the normal interoceptive monitoring taking place while awake. Bodily self-sampling functions to fine-tune the brain's body map, while normal interoceptive monitoring coordinates bodily action (among other things). Exteroceptive sensory processing, while dreaming, functions not to map the external world around us, but to simulate a world matching the modeled body and its actions. If this proposal is correct, then the *normal* functioning of

our sensory systems while asleep and dreaming results in unreliable, illusory, and distorted experiences which don't seem placed to confer justification on beliefs.

Dickie tries to keep her account neutral on the nature of justification (Dickie 2015a, 15). At first she glosses justification in terms of luck (2015a, 43): “if a subject’s belief is justified, the subject will be unlucky if the belief is not true and not merely lucky if it is”. She precisifies this in terms of counterexamples: “a route to belief is justification-conferring iff it eliminates some reasonably wide range of excluded circumstances”, i.e. circumstances in which the belief is false, where a route to belief eliminates a circumstance iff “given that the belief is formed by this route, [the circumstance] cannot be actual” (2015a, 46).

In the end (2015a, 44) she endorses the following principle about justification (and rationality): (Claim A:) “Justification that secures the rationality of a belief eliminates every rationally relevant circumstance where the belief is not true.” From this claim and the further claim (Claim B:) that “If S’s belief that $\langle \alpha \text{ is } \Phi \rangle$ is about an object, the belief is true iff that object is Φ ”, she infers (Claim C:) that “If S’s $\langle \alpha \text{ is } \Phi \rangle$ belief is about o , justification that secures the rationality of the belief eliminates every rationally relevant circumstance where o is not Φ ” (2015a, 48).²³ She takes this argument as demonstration of the left-to-right direction of the following biconditional, which encodes the core thesis of her view: “A body of *ordinary* $\langle \alpha \rangle$ beliefs is about o iff its proprietary means of justification converges on o , so that, for all $\langle \Phi \rangle$, if S has proprietary rationality-securing justification for the belief that $\langle \alpha \text{ is } \Phi \rangle$, this justification eliminates every rationally relevant circumstance where o is not Φ ” (2015a, 57).

What we have is (roughly put) the view that a body of beliefs (a mental file) is about an object iff its proprietary route to justification converges on that object, where this route to justification eliminates every rationally relevant circumstance in which the object fails to have the relevant property. My focus is just on the left-to-right direction, which for clarity I’ll shorten further to: *a mental file is about an object only if its proprietary route to justification eliminates every rationally relevant circumstance in which the object fails to have the relevant property*. This gloss is faithful to Dickie, as she herself says (2015a, 57): “an aboutness-fixing relation does its aboutness-

²³I have labeled these three claims as Claim A, B, and C because they will be important shortly.

fixing work by putting in place a route to justification for a body of $\langle \alpha \rangle$ beliefs such that if the subject *were* to form $\langle \alpha \text{ is } \Phi \rangle$ beliefs justified by this means, these beliefs *would* tend to get *o*'s properties right." Dickie does allow that defeaters might get in the way of justification (2015a, 58), but the idea is that a reference-fixing relation does in fact put in place a route to justification which would provide justification were it not for those defeaters.

While Dickie does provide an argument for the left-to-right direction of her view, this argument is the inference of what I above labeled Claim C from Claims A and B. I agree with Dickie that Claim C is true (and follows from A and B); but the issue is that Claim C is not the same as this more robust claim (on p. 57) about how "an aboutness-fixing relation does its aboutness-fixing work" or even the thinner claim that reference-fixing requires a route which (absent defeaters) provides justification. Claim C, if it's to be read in a way that logically follows from A and B, merely says that if a belief is about an object, then *if* there is justification securing the rationality of the belief, this justification eliminates every rationally relevant circumstance where the object lacks the relevant property.²⁴

Setting aside all the technical details, what we're left with is the question: *is it true that a demonstrative thought is about an object only if the link exploited to think that thought provides a route to justification which, absent defeaters, eliminates every rationally relevant circumstance in which the object fails to have the property attributed to it by the thought?* I take it that Dickie wants to answer "yes", but I think the example of stimulus-incorporating dream experience shows the answer is "no". Conscious attention to a stimulus incorporated into a dream experience opens up a perceptual link to that stimulus which allows for demonstrative selection. But this attention does not constitute a belief-forming route which would, absent defeaters, provide justification or eliminate every rationally relevant circumstance.

This claim is plausible based on the unstable, bizarre, and generally unreliable nature of dream experiences and the way they leave us cut off from the world, but (as noted before) Windt's view, if correct, gives us more insight. As she proposes, the sensory channel which affords the perceptual link in dreams does not function in a way that's sensitive to the actual way things are out in the world. In the case of exteroceptive senses this channel is highly attenuated and when working at all funnels input into a body-centered sim-

²⁴A moment of careful inspection of the three claims, as printed above, should convince the reader that this is true.

ulation. In the case of interoceptive senses this channel facilitates bodily self-sampling instead of (as usual) sensorimotor coordination. So, attention to what's presented in dream experience is a belief forming route which, in the normal course of its operation, gets things very wrong; even absent defeaters, one should not trust it. In fact, finding out that you are connected to a stimulus through this sensory channel—e.g., by realizing that you're dreaming—suffices to discount any beliefs you may have formed based on it.

In her discussion of demonstrative thought (2015a, chapter 4), Dickie spends substantial time arguing that perceptual links do provide justification of the sort required by her account. If my counterexample works, where does this argument go wrong? The crux of her argument rests on the following empirical claim:

The ordinary pathway to formation of ⟨That is Φ ⟩ beliefs is uptake from the property-level appearances delivered by the kind of property-from-feature calculations that our visual systems perform within the attentional spotlight. If, but only if, the attended object is referentially basic, ⟨That is Φ ⟩ beliefs formed by this pathway will reliably match what it is like. (Dickie 2015a, 121)

The property-from-feature calculations mentioned here are those performed by the visual system as it tracks constant object features through changing proximal stimulation and which explain why you see (e.g.) a spinning penny as a constant circular shape despite a changing array of ellipses projecting onto your retina. Given how Dickie uses the term, it's definitional these calculations will be reliable when tracking *referentially basic* objects in normal circumstances (2015a, 120).

Although Dickie integrates this claim into a broader argument which establishes that perceptual links afford justification on her preferred account of justification, this claim underpins the “structural core around which any candidate account of perceptual demonstrative aboutness-fixing should be built: an attentional perceptual link with a referentially basic object enables perceptual demonstrative thought about it in virtue of the fact that, where the attended object is referentially basic, ⟨That is Φ ⟩ beliefs justified by the *standard route* will reliably get its properties right” (2015a, 130). By the ‘standard route’, Dickie means “uptake from the property-level appearances delivered by the kind of property-from-feature calculations that our visual systems perform within the attentional spotlight” (2015a, 121). The problem is that while the standard route to belief formation *is* reliable when it's the

visual system doing these attention-aided property-from-feature calculations, the standard route isn't always reliable when it's other sensory channels at work. For example, one's own body is a referentially basic object (or so seems plausible) and in dreams we can form beliefs based on uptake from the property-level appearances returned by attention to the interoceptive sensory channel engaged in bodily self-sampling, but this sensory channel is not reliable. Hence, Dickie overgeneralizes from the case of vision; the standard route, when it involves dream bodily self-sampling and other attenuated or simulation-driven sensory processes, isn't reliable.

9 Concluding Remarks

I've argued that the specific epistemic necessary conditions placed on demonstrative selection by proponents of the epistemic approach do not hold. These proponents are Evans, Recanati, and Dickie, although much of their picture is accepted in bits and pieces by other philosophers (such as Jeshion (2010) and Smithies (2019)). This approach has it that demonstrative selection requires an epistemically robust information link and that perception enables demonstrative selection by providing such a link. In addition, Evans proposes that demonstrative selection requires knowledge of a target's position in public space, while Dickie says that it requires a link which provides justification. I've argued that in the cases of multimodal misbinding and stimulus-incorporating dream experience we still have a perceptual link which enables demonstrative selection, but between these two cases we have a counterexample to each of these requirements.

It might be objected that my argument rests on intuitions about the possibility of demonstrative selection in these two cases, and that such intuitions don't carry much argumentative weight. But first, my case isn't purely intuitive. I've given a reason for accepting the possibility of demonstrative thought in these examples: both involve genuine perceptual links. Second, my opponents themselves make use of intuitive cases to motivate their accounts (Evans 1982, 164–67), acknowledge the role these intuitions must inevitably play in theorizing (Dickie 2015a, 19), and acknowledge that their accounts should be at least somewhat responsive to these intuitions (Dickie 2015a, 22). Finally, I have not merely given counterexamples. I've outlined my opponents accounts in detail and identified the precise steps at which they go wrong.

Acknowledgements

This paper owes a great deal to the two anonymous referees from this journal, whose sharp comments greatly improved both the framing of this paper and its central arguments. This paper grew out of a draft of a much different paper, which itself started as chapter 3 of my dissertation. In that long process I benefited greatly from input by Casey O’Callaghan, Charles Siewert, Nico Orlandi, Indrek Reiland, Richard Grandy, Dan Burnston, Alex Morgan, Mohan Matthen, and several other sets of referees. Casey, Indrek, Alex, and Mohan deserve special mention, as they provided generous line-by-line comments on earlier drafts and spent substantial time discussing with me fundamental issues which shaped this paper into its current form. This paper was written during my postdoc under Mohan Matthen at the University of Toronto. I thank Mohan and the University for their generous support.

References

- Allport, A. (2011). Attention and integration. In C. Mole, D. Smithies, and W. Wu (Eds.), *Attention: Philosophical & Psychological Essays*, pp. 24–59. Oxford: Oxford University Press.
- Bach, K. (2010). Getting a thing into a thought. In R. Jeshion (Ed.), *New Essays on Singular Thought*, pp. 39–63. Oxford: Oxford University Press.
- Barkasi, M. (2019). The role of experience in demonstrative thought. *Mind and Language Feb 3 (online first)*, 1–19.
- Blumberg, M. S. (2010). Beyond dreams: Do sleep-related movements contribute to brain development? *Frontiers in Neurology 1(140)*, 1–10.
- Blumberg, M. S. (2015). Developing sensorimotor systems in our sleep. *Current Directions in Psychological Science 24(1)*, 32–37.
- Blumberg, M. S., C. M. Coleman, A. I. Gerth, and B. McMurray (2013). Spatiotemporal structure of REM sleep twitching reveals developmental origins of motor synergies. *Current Biology 23(21)*, 2100–2109.
- Burge, T. (2010). *Origins of Objectivity*. Oxford: Oxford University Press.
- Campbell, J. (2002). *Reference and Consciousness*. Oxford: Oxford University Press.
- Campbell, J. (2004). Reference as attention. *Philosophical Studies 120*, 265–276.

- Carrasco, M. (2011). Visual attention: The past 25 years. *Vision Research* 51(13), 1484–1525.
- Cheyne, J. A. (2003). Sleep paralysis and the structure of waking-nightmare hallucinations. *Dreaming* 13(3), 163–179.
- Clark, A. (2006). Attention & inscrutability: A commentary on John Campbell, “reference and consciousness” for the Pacific APA Meeting, Pasadena, California, 2004. *Philosophical Studies* 127, 167–193.
- Clark, A. (2012). Dreaming the whole cat: Generative models, predictive processing, and the enactivist conception of perceptual experience. *Mind* 121(483), 753–771.
- Clark, A. (2013). Whatever next? predictive brains, situated agents, and the future of cognitive science. *Behavioral and Brain Sciences* 36, 181–253.
- Dement, W. and E. A. Wolpert (1958). The relation of eye movements, body motility, and external stimuli to dream content. *Journal of Experimental Psychology* 55(6), 543–553.
- Dickie, I. (2010). We are acquainted with ordinary things. In R. Jeshion (Ed.), *New Essays on Singular Thought*, pp. 213–245. Oxford: Oxford University Press.
- Dickie, I. (2011). Visual attention fixes demonstrative reference by eliminating referential luck. In C. Mole, D. Smithies, and W. Wu (Eds.), *Attention: Philosophical & Psychological Essays*, pp. 292–322. Oxford: Oxford University Press.
- Dickie, I. (2015a). *Fixing Reference*. Oxford: Oxford University Press.
- Dickie, I. (2015b). Perception and demonstratives. In M. Matthen (Ed.), *Oxford Handbook of Philosophy of Perception*, pp. 833–852. Oxford: Oxford University Press.
- Dretske, F. (1988). *Explaining Behavior: Reasons in a World of Causes*. Cambridge: The MIT Press.
- Evans, G. (1982). *Varieties of Reference*. Oxford: Oxford University Press.
- Hobson, J. A. (2009). REM sleep and dreaming: Towards a theory of proto-consciousness. *Nature Reviews Neuroscience* 10(11), 803–813.
- Hobson, J. A. and K. J. Friston (2014). Consciousness, dreams, and inference: The cartesian theatre revisited. *Journal of Consciousness Studies* 21(1-2), 6–32.
- Hobson, J. A., C. C. H. Hong, and K. J. Friston (2014). Virtual reality and consciousness inference in dreaming. *Frontiers in Psychology* 5(1133), 1–18.

- Hoff, H. (1929). Zusammenhang von Vestibularfunktion, Schlafstellung and Traumleben. *European Neurology* 71(5-6), 366–372.
- Hoff, H. and O. Pötzl (1937). Über die labyrinthären Beziehungen von Flugsensationen und Flugträumen. *European Neurology* 97(4), 193–211.
- Jeshion, R. (2010). Singular thought: Acquaintance, semantic instrumentalism, and cognitivism. In R. Jeshion (Ed.), *New Essays on Singular Thought*, pp. 105–140. Oxford: Oxford University Press.
- Johnston, M. (2011). On a neglected epistemic virtue. *Philosophical Issues* 21(1), 165–218.
- Kahneman, D. and A. Treisman (1984). Changing views of attention and automaticity. In R. Parasuraman and D. A. Davies (Eds.), *Varieties of Attention*, pp. 29–61. New York: Academic Press.
- Kahneman, D., A. Treisman, and B. Gibbs (1992). The reviewing of object files: Object-specific integration of information. *Cognitive Psychology* 24, 175–219.
- Leslie, K. and R. Ogilvie (1996). Vestibular dreams: The effect of rocking on dream mentation. *Dreaming* 6(1), 1–16.
- Levine, J. (2010). Demonstrative thought. *Mind and Language* 25, 169–195.
- Martin, M. G. F. (2002). Particular thoughts and singular thought. In A. O’Hear (Ed.), *Logic, Thought, and Language*, Volume 51 of *Royal Institute of Philosophy Supplements*, pp. 173–214. Cambridge: Cambridge University Press.
- Matthen, M. (2005). *Seeing, Doing, and Knowing: A Philosophical Theory of Sense Perception*. Oxford: Oxford University Press.
- Matthen, M. (2012). Visual demonstratives. In A. Raftopoulos and P. Machamer (Eds.), *Perception, Realism and the Problem of Reference*, pp. 43–67. Cambridge: Cambridge University Press.
- McDowell, J. (1994). *Mind and World*. Cambridge: Harvard University Press.
- Millikan, R. G. (2012). Are there mental indexicals and demonstratives? *Philosophical Perspectives* 26(1), 217–234.
- Nielsen, T. A. (1992). A self-observational study of spontaneous hypnagogic imagery using the upright napping procedure. *Imagination, Cognition and Personality* 11(4), 353–366.
- Nielsen, T. A. (1993). Changes in the kinesthetic content of dreams following somatosensory stimulation of leg muscles during REM sleep. *Dreaming* 3(2), 99–113.

- Nielsen, T. A., L. Ouellet, and A. Zadra (1995). Pressure stimulation during REM sleep alters dream limb activity and body bizarreness. *Sleep Research* 24, 134.
- O’Callaghan, C. (2007). *Sounds: A Philosophical Theory*. Oxford: Oxford University Press.
- O’Callaghan, C. (2015). The multisensory character of perception. *Journal of Philosophy* 112(10), 551–569.
- O’Callaghan, C. (2016). Objects for multisensory perception. *Philosophical Studies* 173, 1269–1289.
- O’Callaghan, C. (2019). *A Multisensory Philosophy of Perception*. Oxford: Oxford University Press.
- Peacocke, C. (1983). *Sense and Content: Experience, Thought, and Their Relations*. Oxford: Clarendon Press.
- Pylyshyn, Z. W. (2007). *Things and Places: How the Mind Connects with the World*. Cambridge: The MIT Press.
- Raftopoulos, A. and V. Müller (2006). Nonconceptual demonstrative reference. *Philosophy and Phenomenological Research* 72(2), 251–285.
- Recanati, F. (2012). *Mental Files*. Oxford: Oxford University Press.
- Russell, B. (1997/1912). *The Problems of Philosophy*. Oxford: Oxford University Press.
- Sauvageau, A., T. A. Nielsen, and J. Montplaisir (1998). Effects of somatosensory stimulation on dream content in gymnasts and control participants: Evidence of vestibulomotor adaptation in REM sleep. *Dreaming* 8(2), 125–134.
- Schellenberg, S. (2019). Perceptual consciousness as a mental activity. *Nous* 53(1), 114–133.
- Shannon, C. (1948). A mathematical theory of communication. *Bell System Technical Journal* 27, 379–423, 623–656.
- Siegel, S. (2006). How does visual phenomenology constrain object-seeing? *Australasian Journal of Philosophy* 84(3), 429–441.
- Siegel, S. (2010). *The Contents of Visual Experience*. Oxford: Oxford University Press.
- Siewert, C. (1998). *The Significance of Consciousness*. Princeton: Princeton University Press.
- Smithies, D. (2011). What is the role of consciousness in demonstrative thought? *Journal of Philosophy* 108(1), 5–34.
- Smithies, D. (2019). *The Epistemic Role of Consciousness*. Oxford: Oxford

University Press.

- Treisman, A. (1998). Feature binding, attention and object perception. *Philosophical Transactions of the Royal Society B: Biological Sciences* 353, 1295–1306.
- Windt, J. M. (2018). Predictive brains, dreaming selves, sleeping bodies: How the analysis of dream movement can inform a theory of self- and world-simulation in dreams. *Synthese* 195, 2577–2625.
- Wu, W. (2011). Attention as selection for action. In C. Mole, D. Smithies, and W. Wu (Eds.), *Attention: Philosophical & Psychological Essays*, pp. 97–116. Oxford: Oxford University Press.
- Zmigrod, S., M. Spapé, and B. Hommel (2009). Intermodal event files: integrating features across vision, audition, taction, and action. *Psychological Research* 73(5), 674–684.