

Thought and Talk in a Generous World

Alexander Sandgren (Umeå University)

forthcoming in *Ergo: An Open Access Journal of Philosophy*

Abstract

The problem of the many seems to problematize the platitude that we can think about particular things in the world. How is it that, given how very many cat-like candidates there are, we often manage to think and talk about a particular cat? I argue that this challenge stems from an under-examined assumption about the relationship between metaphysics and intentionality. I explore and develop a way of characterizing what it is to think and talk about the world, according to which an abundant ontology poses no obstacle to our ability to think and talk about particular things.

1 The Needle and the Haystack

A hero is summoned to find a single needle in a haystack. Confronted with a heap of nearly indistinguishable prickly pieces, her task looks impossible. After all, it is unlikely that the world will be much help (the needle isn't glistening conveniently on the surface) and her own ingenuity is no match for the size of the job.

According to certain common views of the relationship between representation and the world, ordinary attempts to think and talk about the world seem to present us with the same challenge faced by our hapless hero. We often think and talk about particular objects, and much seems to depend on *what* we're thinking and talking about. Meanwhile, there are good reasons for thinking that the world is *generous*: that there are ever so many mereological composites, spatiotemporal coincidents, properties, contents, and so on. Against the backdrop of worldly generosity, if we think that in order to think and talk about particular objects in the world we must somehow select a single object from the many then our task is like the task of pulling a single needle from a haystack. So this generosity poses a challenge: is the task of ordinary thought and talk getting to be about particular things akin to pulling a needle from a haystack, and if so, how are we so often successful in thinking about

particular objects? How do we manage to think and talk about particular things, when the world supplies so many candidates as potential intentional targets?

In what follows, I examine some familiar strategies for answering these questions, and defend what I take to be an underappreciated route. The two main contributions are i) parts of the negative discussion in which I raise novel problems for familiar approaches to the problem of the many and ii) the use of considerations based on the problem of the many to motivate and elaborate an approach to aboutness, recently outlined in my ‘Turning Aboutness About’ (Sandgren 2021). I will answer many questions left open in that paper and propose an account of how aboutness relates to questions about quantification over intentional objects, truth, and counting.

I will focus primarily on variants of *the problem of the many* for material objects, but much of what I will say may generalize to other puzzles about thought and talk in the face of generosity. I will often write in terms of ‘talk’ rather than ‘thought and talk’, but I mean most of my remarks to apply widely to other puzzles concerning other kinds of representations including intentional attitudes.

2 Generosity and Problems of the Many

Standard presentations of the problem of the many concern settings where mereological composites are abundant.¹ For example,

Fuzzy. Ponsonby the cat is sitting on the mat, apparently alone. As with most cats, there are a *lot* of hairs wherever she is: $h1, h2 \dots h1001$. But given *mereological universalism*, there are many material things on the mat: something which has all but $h1$ as parts; something which has all but $h2$ as parts, \dots and so on.

Here is an observation from Geach (1980) via Lewis (1993), on the observation of *questionable parts*:

‘Once noticed, we can see that it is everywhere, for all things are swarms of particles. There are always outlying particles, questionably parts of the thing, not definitely included and not definitely not included. So there are always many aggregates, differing by a little bit here and a little bit there, with equal claim to be the thing. We have many things or we have none, but anyway not the one thing we thought we had. That is absurd.’

¹There is a large and varied literature on the problem of the many; see especially Unger (1980), Geach (1980), and Lewis (1993). It is a further interesting question how the issues I discuss here arise for mereological eliminativists; see McGrath (2005) and Rettler (forthcoming) for a related discussion.

Even given only the observation that mereological composites are abundant, together with an understandable reluctance to split hairs about the mereological makeup of cats, we find ourselves in dire straits. But the trouble isn't (or needn't be) just a matter of mereology and 'questionable parts'.

Other observations about material abundance – like temporal and modal plenitude – threaten to make matters even worse. Each entails, respectively, an abundance of *spatially* coincident objects and *spatiotemporally* coincident objects. Perhaps the most straightforward formulation of plenitude comes from Hawthorne (2006): roughly, there is an object corresponding to every function from worlds to filled regions of spacetime. The details of the puzzling cases I appeal to here will not turn on the details of the formulation of plenitude. See Fairchild (2019; 2020) for more on the upshots of varieties of formulations of plenitude. Lewis' discussion of the problem of the many invokes all three dimensions of abundance, though given his background metaphysics, they're all in some sense a matter of *questionable parts*. Consider:

Slow Change. Ponsonby the cat is sitting on the mat, apparently alone. Like most cats, she sheds mercilessly. Suppose that this spring, 1001 hairs will be lost: $h_1, h_2 \dots h_{1001}$. Given *temporal plenitude*, there are therefore very many things on the mat: something which won't survive losing h_1 ; something which won't survive losing h_2 , ... and so on.

Possible Change. Ponsonby the cat is sitting on the mat, apparently alone. Like most cats, she sheds mercilessly. This Spring, 1001 hairs will be lost: $h_1, h_2 \dots h_{1001}$. But winter always comes, so thankfully, 1001 hairs will stick around, too: $h_{1002}, h_{1003}, \dots h_{2003}$. Given *modal plenitude*, there are therefore very many things on the mat: something which in fact survives, but wouldn't have survived losing h_{1002} ; something which wouldn't have survived losing h_{1003} , and so on.²

These cases arise from commitments to the abundance of the material world which are naturally (though certainly not inevitably) seen as a package deal. Taken together, the observations above deliver a *very* full mat.

All of this is to say: the world is generous enough to make our lives quite difficult, if indeed abundance makes our lives difficult. But what exactly does that

²Caution here: as I noted above, we aren't *guaranteed* cases like this by every version of plenitude – for example, not all versions of plenitude will entail that there are mereologically complex objects (like cats!). See Fairchild (2019) and Korman (2015). I am also following Lewis in my choice of example, but you might think that cats aren't the most illustrative case: plausibly nothing sufficiently cat-like could be destroyed by hair loss. If it helps to avoid distraction, perhaps consider cases involving slow change – for example, cats whose organic parts are gradually replaced by robotic parts.

difficulty consist in? As I see it, material abundance makes it hard to answer a family of related questions:

COUNTING CATS. How many cats are on the mat? (That is, how many of these things are *cats*?)

HITTING TARGETS. How, in saying things like ‘the cat on the mat has black fur’, have we managed to successfully talk about a particular cat, given how very many cat-candidates there are?³

ACHIEVING TRUTH. How do we manage to say true things in a generous world? (for example ‘The cat on the mat has black fur.’)

In the literature on the problem of the many, COUNTING CATS and ACHIEVING TRUTH seem to have received the bulk of the attention. However, part of what I aim to do here is make clear how previous proposals may provide satisfying answers to both of these questions without thereby saying anything very satisfying at all about HITTING TARGETS. Here I will be mainly concerned with HITTING TARGETS, a question at the center of what Openshaw (2020) aptly calls ‘the intentional problem of the many’. Despite this focus, the relationship between HITTING TARGETS and the other two questions will be relevant when I discuss the positive proposal in section 5.

Before moving on, I want to bring the reader’s attention to a feature of the literature on these questions. For that, I’ll help myself to some rough imagery: let’s call the ultimate inventory of what there is (given perfect ontological accounting) ‘the list’. The challenges associated with generosity inherit much of their force from the assumption that the intentional features of our thought and talk are beholden to the list and from the idea that someone’s talking about, for example, a particular cat is explained by there being some particular object selected from the list as *the* intentional target. This idea is natural for those who adopt what I (2021) call a *target-first approach* to aboutness, so called because of the focus on intentional targets (the entries on the list) as explaining aboutness. ‘Target-firsters’ hold that adequate theories of aboutness – and adequate answers to HITTING TARGETS – partly explain

³This ‘aboutness’ is importantly different from what has been discussed under that name by Yablo (2014), Fine (2014; 2016), and others. The notion of representational aboutness at stake in Yablo (2014), for example, is tied up with what makes (or could make) a representation true. However, following Lewis (1981; 1986: 33) I suggest that questions of aboutness can at least in principle be detached from truth conditions. We’ll see in what follows that there is reason to separate these questions so as to leave space for views on which what a representation is about (i.e. what the representation is a representation of) may not supervene on its truth conditions. So, for example: Pierre’s belief that *London is pretty* might be true just in case there is a Londonish thing that is pretty – but his belief could be *about* London all the same. For a recent discussion of this point and an exploration of how aboutness and truth conditions might come apart, see Sandgren (2019b: section 4).

aboutness in terms of the list and entries on the list. For example, a target-firster will claim that a representation, say a name, has the intentional features it has partly *because* there is an entry or set of entries on the list having certain features (e.g. they satisfy some description, or they stand in some causal relation, or they are the object with which the representer is acquainted). The candidate intentional targets of the representations (the entries on the list) are crucial for explaining the intentionality of representations. In this sense these proposals are 'target-first'.

The target-first approach is familiar to the realists among us: the world supplies the catalog of what there is, and theories are built around and constrained by the catalog. What we can think and talk about is constrained by what there is, or so the idea goes. So insofar as we successfully engage the world with our thought and talk, we do so by 'connecting up' with entries on the list. But although lovers of the list seem to have a monopoly on realism, much of the goal of what follows is to complicate that connection. As I see it, there is underappreciated space for a realist account according to which we manage to talk and think about the world and things in the world – to talk and think about cats and clouds and churches – without being beholden to the list.

3 Two Approaches

3.1 We Manage to Select One of the Many Without Help: 'the Artificial Glow'

Perhaps we manage to 'pick one from the many' thanks to the richness of the linguistic conventions and intentions. After all, the intention to talk about certain kinds of things may narrow the field (e.g. the intention to talk about a cat may rule out some of the many – the non-cats). There may also be interpersonal conventions that impose their own constraints on what we are talking about. Our intentions and conventions are extremely rich: rich enough, one might think, to single out one object from the many. On this proposal there is one object among the many endowed with an 'artificial glow' as a result of our individual or collective practices and intentions.

But many of the differences among the many are such that our conventions and intentions do not suffice to *uniquely* select one from the rest. For example, there are objects that differ only with respect to the inclusion of one electron. But plausibly our conventions do not distinguish between a cat and a cat plus one extra electron. Plausibly our conventions associated with cats and 'cat' are just not sensitive to the inclusion of the extra electron. And certainly the conventions of humans that have never heard of electrons will not be sensitive to such a difference. But these are

different objects with, it seems, equal right to count as cats!⁴

Those who would claim that our conventions and intentions single out one of the many seem not to have taken the scale of worldly generosity seriously enough: the world is abundant along too many dimensions, with too much fine-grained variation, for our practices to tame. Our intentions and conventions are simply too blunt to make the required distinctions.

3.2 The World Helps Us: 'the Natural Glow'

Maybe our conventions and intentions do not have to do all of the work. Recently, there has been significant interest in the question of whether there are constraints on the interpretation of our thought and talk that do not depend wholly on us, our psychology, or conventions.⁵

Perhaps the very many objects in the vicinity are not all on equal footing: maybe some are more eligible as targets for our thought and talk. According to this story, our conventions and intentions together with worldly 'eligibility' suffice to select one from the many.

The most familiar application of the worldly eligibility story is to properties: we say that some properties are more 'joint carving' or 'natural' than others, and thus more eligible as targets for thought and talk. For instance, the property of being tired or grey is plausibly less natural than the property of being grey, the set of grey things are more unified, in some crucial sense, than the set of tired or grey things. Lewis (1983, 1984) often speaks of comparative eligibility relations between properties but also contends that objects can be relatively eligible referents in virtue of instantiating relatively natural properties (Lewis 1983: 372). But, as Lewis notes, it isn't clear how those resources could help with the problem of the many we are discussing, 'nature is gradual, no handy joint in nature picks out one of the [objects] from all the rest.' (Lewis 1993: 172, n.6)

As a first pass, there are two ways one might try to use property eligibility to help with the intentional problem of the many: ranking objects by their descriptions in terms of perfectly natural properties, or by the comparative naturalness of their complete qualitative descriptions.

⁴You might think that the answer to counting cats is easy: there is only one *cat* among the many. Thus, to answer HITTING TARGETS, all we need is *the intention to refer to a cat* to ensure that we select one from the many (Lowe 1982; 1995). But this seems to push the bump under the rug: the proponent of such a proposal will have to say that either (i) our conventions guarantee that only one thing on the mat will satisfy the predicate 'is a cat such that we are thinking about it' or that (ii) one thing is metaphysically privileged with respect to catness. The latter idea is a version of the proposal discussed in section 3.2. below and the former runs into the worries to be discussed in the rest of this section.

⁵See, for instance, Lewis (1983; 1984), Sider (2009; 2011: 23-35), and Williams (2007; 2015; 2020).

Once again, an obstacle to both strategies is that, having taken plenitudinous metaphysics seriously, we'll find that many of the objects the world offers will differ only with respect to very unnatural properties. What this means for the first strategy is that if we attempt to pick out an individual with a description that mentions only perfectly natural properties, there will likely still be more than one object that fits that description. For example, if two candidate targets have all the same perfectly natural properties but differ with respect to some relatively unnatural property (for example the property of being closer to Betelgeuse) they will still be distinct objects, and yet both will be equally good candidates to be picked out by the description in terms of perfectly natural properties.

The second strategy seems similarly likely to yield ties. Consider again the mereological case: which has a more natural complete qualitative description, the object which doesn't have *h1* as a part but has *h2*, or the object which has *h2* but not *h1*? Of course, these objects will differ with respect to other very unnatural properties, but it seems unduly optimistic to hope that this will suffice to break these ties.⁶

But perhaps objects get their eligibility directly and not via the relative eligibility or naturalness of their properties. We might hope to find a solution to the present difficulties by appeal to a new notion of comparative eligibility for material objects. Perhaps some objects are like Toner's (2008) emergent simple objects with no substantial parts and are for that reason eligible as intentional targets, or perhaps macro-level objects like cats are fundamental (Bernstein 2021) and therefore relatively eligible.⁷ What should we think of this move? Like the indirect approach, the object eligibility picture invokes a number of heavy duty metaphysical commitments. Some of the many are supposed to be special in some way that means that they are more eligible to intentional objects than other objects.⁸

For that reason it is worth exploring less metaphysically costly alternatives. Although it might turn out that, at the end of the day, the price is worth it to avoid even more unpalatable alternatives, we do better to exhaust the space of

⁶One further option is to say that comparative eligibility relations stand between *identity properties* of objects; but this version of property eligibility appears to be object eligibility in different clothes.

⁷We often talk about more and less 'natural' objects: for example, perhaps electrons are more natural than cats (e.g. because electrons appear are more fundamental than cats). But it is not clear that this kind of move will help here. It seems implausible that any of the composites in the vicinity of the cat is more fundamental (and thus relevantly more natural) than all or most of the others. See Brown (2016) for a related discussion of attempts to extend a notion of 'naturalness' to objects.

⁸Another suggestion is that, in some sense, the privileged macro-level level object is the *only* object in the vicinity. The idea would be that there is one cat, perhaps it has an Aristotelian substantive form, and all the 'many' are not *really* substances or, at least, candidates to be intentional objects. This move does get one out of some of the problems of the many, but that is not really surprising, since it essentially involves the claim that the world is not really generous. But since my focus is on how intentionality works *given that the world is generous*, I will set that line of discussion aside for now.

options before signing the check. Relatedly, but more controversially, this may seem somewhat suspect as a dialectical move: faced with the many, we are reassured that only a privileged few will have the convenient but mysterious glow of worldly eligibility. At least as an answer to HITTING TARGETS, this doesn't strike me as a very satisfying explanation.

4 One does not need to be selected from the many...

Each of the approaches above has in common the thought that a unique candidate must *glow* for our thought and talk to latch on to the world in a respectable way. In this section, I explore another kind of strategy for approaching the puzzle: perhaps thinking and talking about the world doesn't require selecting one from the many. Maybe, in some sense, our thought and talk about the world is just more coarse-grained than the world itself. It is illustrative here to draw an analogy with *seeing* in a generous world. There are a host of quite delicate perceptual issues about which material objects we *see* when we look at, for example, a cat on a mat, given plenitude.⁹ But however these are to be resolved, it seems clear that it would be surprising if it turned out that we *only* see exactly one 'fine-grained' material object. Rather, it is much more tempting to look for a solution that respects the plausible thought that our perceptual systems just don't carve the world that finely. Below, I consider two ways of implementing this idea for the case of aboutness.

4.1 ...because a class of objects will do.

Perhaps we succeed in talking about the world not by successfully singling out one object from the many, but instead by carving out a *class* of objects. At first blush, this strategy promises to capture the tempting idea that our thought and talk are more coarse-grained than the world itself, while nonetheless preserving a central role for the list in explaining and constraining our thought and talk. Thus this strategy is an attempt to deny that we need to single out one object while hanging on to the general spirit of target-first approaches.

The paradigmatic instance of this approach employs the method of supervaluations. Take the claim that 'the cat on the mat has black fur'. The core supervaluationist idea is that there is a class of objects in the vicinity that are *good enough* candidate referents for our cat-talk: each such object is understood to be something that we – depending on further semantic choices – could straightforwardly refer to. Call this 'the class of acceptable referents' (CAR).¹⁰ The supervaluationist can

⁹Some arise already with universalism: when I see Ponsonby, do I see *each* of the bizarre fusions that has Ponsonby as a part, just as I see you when I see only your right arm?

¹⁰On some versions of the plenitudinous views at issue here, there will be too many candidates to

then say that ‘the cat on the mat has black fur’ is *super-true* just in case all of the members of the relevant CAR have black fur. This is naturally paired with a further claim about our practices: we often aim for super-truth, rather than truth *simpliciter*. Thus, despite the complications posed by worldly generosity, we may be entitled to say that the cat on the mat has black fur.

Supervaluationist tools promise to help us answer ACHIEVING TRUTH and COUNTING CATS. Roughly, the supervaluationist strategy with respect to the former is to say that we manage to speak truly in the face of generosity by managing to speak super-truly. The strategy for COUNTING CATS is a bit more delicate, but stems from the supervaluationist observation that ‘there is exactly one cat on the mat’ is often super-true. If we attend to admissible precisifications of the predicate ‘is a cat’, the thought goes, we’ll notice that ‘there is exactly one cat on the mat’ is true on every admissible precisification. But even so, it isn’t obvious that this toolkit will thereby help us answer HITTING TARGETS. As Openshaw (2020) argues convincingly, *vanilla* supervaluationism leaves the intentional problem of the many unsolved. Broadly speaking, it isn’t clear how to connect either the method of supervaluations or the notion of super-truth to questions of *aboutness*.

There are three salient ways of trying to answer HITTING TARGETS by appealing to distinctively supervaluationist resources. Some options are to claim that i) we talk about all of the members of the CAR, ii) there is a particular member of the CAR such that we determinedly talk about it, or iii) none of the members of the CAR are such that we talk about it.

There are two disambiguations of the first option – a *distributive reading* and a *lumping reading*. On the distributive reading, in talking about ‘the cat on the mat’, we talk about *each* member of the CAR. This is roughly the route taken by Openshaw (2020), Merlo (2017), and Hawthorne and McGonigal (2008) who argue that paradigmatic expressions of singular thoughts – like ‘the cat is on the mat’ – express *many* contents (of perhaps many thoughts), each singularly about ‘the one cat’ on the mat.

Imagine you say that the cat in the kitchen is grey but unbeknownst to you there are, as we would usually say, two cats in the kitchen. In cases like this you have failed to talk about about a particular cat because there are too many in the kitchen. But there is nothing in the proposal we are currently discussing that stops us from simply claiming that there are multiple referents in this case, just as there are in cases in which we would ordinarily say there is one cat in the kitchen. Those who take the distributed reading of the ‘thinking about all of them’ option will still need a way of distinguishing the cases in which the many lead to a failure to think about a particular cat from those in which the multiplicity of targets is benign. Such an explanation would be, in effect, a solution to the intentional problem of the many,

form a set, so I’ll stick to the language of classes.

so we are back where we started.

It isn't better to retreat to a lumping reading of i), and answer that we are talking about all of the candidates taken together – and so, in a sense, we are talking about a single thing (namely, a class). But consider that 'Queen played at Live Aid' is about a band not a class of bands. Our conviction that we are talking about one cat (not a class of cats) is at least as strong as our conviction that we're talking about one thing. Queen is not made of bands, it is made up of people. But the set of candidate referents on the mat (certainly many of them) are all plausibly cats (they walk like cats, purr like cats...) and it is bizarre to say that we are thinking about a class of cats. In sum: on the distributive reading we are talking about many things, on the lumping reading we are talking about one thing (the class), but on neither are we talking about one cat, so we will need to look elsewhere for an answer to HITTING TARGETS.

As for (ii); in the absence of other constraints, any choice of target from among the CAR would be objectionably arbitrary such that exploring alternative approaches seems worthwhile.¹¹

Another route (iii) is to concede that there is no member of the CAR such that we are talking about it. This is the option that Lewis (1993; 1999: 9) seems to favor. There Lewis is concerned with the status of existential generalizations: it is supertrue that the cat is on the mat, though there is no cat of which it is supertrue that *it* is the cat on the mat. His envisioned objector remarks: 'It's like the old puzzle: I owe you a horse, but there's no horse such that I owe you that horse. This is peculiar.' Lewis (1993: 173) replies: 'So it is. But once you know the reason why, you can learn to accept it.' The analogous move here would be to admit that there is no cat on the mat such that we are talking about it.¹²

Lewis' analogy with horse debts might reassure those concerned about the asymmetry between existentials and their instances with respect to supertruth. But with respect to answering HITTING TARGETS, this response leaves a lot to be desired. There is still a feature of thought and talk to be explained, that is the particularity of aboutness. There is something wrong with claiming that Harriet owes you a horse and Tony owes you *the same horse*, but there is no horse they owe you. Tony's debt is constrained by Harriet's: insofar as they owe you the same horse, there must be some sense in which there is a particular horse they owe, however that particularity is to be understood. Analogously, suppose that both Harriet and Tony are talking about the cat on the mat. There is something wrong with

¹¹A further option is to employ some notion of eligibility alongside the supervaluationist machinery, along the lines of Weatherson (2003). See section 3.2. for concerns about this move.

¹²According to McGee and McLaughlin (2000: 142), supervaluationist tools help us count cats, but do not help us understand aboutness. They take one lesson of the problem of the many to be that we seldom genuinely think about particular objects at all.

claiming that Harriet is talking about the cat and Tony is talking about the same cat, but there is no candidate referent such that it is determinately true (or super-true) that *it* is uniquely what they are both talking about. The lesson here is that the aboutness of thought and talk must often leave room for *talking about the same thing* and the associated particularity, however these features are to be understood; understanding the aboutness of thought and talk in a necessarily indefinite way (as being about ‘a cat’ rather than ‘*the* cat’) misses an important feature of aboutness.

My goal in this section was not to undermine supervenient tools or supervenientism more generally. Instead, I argued that supervenientism will at least have to be supplemented before we can answer HITTING TARGETS and develop an account of aboutness in a generous world. The method of supervenients may give us the resources to answer the other two questions I started with. With this in mind, we can simply invite the supervenientist to supplement their account with the resources I am about to develop.

5 Listless Aboutness

5.1 The list and being singular

The supervenientist moves discussed in the previous section attempt to make sense of aboutness while (at bottom) retaining the machinery of a target-first story; the supervenients are over fine-grained referential interpretations. I argued that although supervenientist tools provide a useful replacement for truth, they don’t thereby provide an adequate account of aboutness. In what follows, I build on a different approach, expanding on some of my recent work (Sandgren 2021), particularly the rejection of target-first approaches to aboutness and the conception of aboutness as being explanatorily prior to intentional objects. One contribution is that here I will analyse how my approach connects to and sheds light on the intentional problem of the many (something that I only briefly touch on in earlier work) and how the intentional problem of the many informs how my overall framework should be fleshed out. I will also make progress on a number of issues I left open leaves open, particularly, quantifying into attitude and speech reports, the relationship between aboutness and truth conditions, and how to individuate intentional objects given one eschews list-based explanatory resources.

The first part of the proposal involves a kind of shifting of the goalposts when it comes to answering HITTING TARGETS. ‘Target-firsters’ understand HITTING TARGETS as a question about how to secure reference to one or more entry on the list. On that way of understanding the problem, aboutness is intimately connected with reference; for example, a name is about whatever entry on the list it refers to. Crucially, paradigmatic ‘target-firsters’ hold that the presence of an intentional

target also *explains* how a representation manages to be about something; the name 'Ponsonby' is about a particular cat *because* it refers to some particular cat on the list. This picture fits most naturally with a Millian view of names and the view that singular terms are able to be used when expressing propositions composed partially by the objects the term refers to. But the target-firster might also have views that are closer to the more traditional Fregean conception of thought and talk. For instance, Chalmers (2002) understands aboutness as being tied to reference, while claiming that there is a kind of content that is non-referential.

By distinguishing between ACHIEVING TRUTH and HITTING TARGETS I opened the door to accounts according to which truth conditions and reference on the one hand, and aboutness on the other can come apart. If one is convinced that the truth conditions of the belief that Ponsonby is on the mat involves Ponsonby (the entry on the list), then the question of what the belief picks out would require an answer before we could assess its truth conditions. But as we have seen, the target-first referential approach (even when it is supplemented with supervaluationist tools) does not seem to deliver a satisfactory account of aboutness in the face of worldly generosity. Indeed Weatherson (2003: 489), drawing on some work from McGee and McLaughlin (2000), argues as much and appeals to the problem of the many as a reason to reject the idea that singular beliefs have characteristically singular-object-involving contents. If one understands content as referential, this option is unavailable. As Weatherson (2003: 489) also suggests, understanding the content of singular beliefs in terms of object-involving propositions is not the only way to make sense of the singularity (or *de re*ness) of those representations. There are a number of other ways we can understand the sense in which representations are about particular objects. But if the aboutness of a representation cannot be explained in terms of the list, what *does* aboutness amount to, and what is it grounded in?

5.2 Directedness

We noted above that Lewis' bullet-biting on the aboutness question was unsatisfying because targeted thought and talk *aren't* like indefinite descriptions. Beliefs are often not about some cat or other, but about some particular cat. Definite descriptions ('the cat in the kitchen') are directed at particular things in a way that indefinite descriptions ('a cat', 'a horse') are not. Russell (1905: 79) characterises definite descriptions as denoting phrases that do not always have a denotation. Even though 'the present queen of France' does not denote anything, it is nonetheless a denoting phrase. The idea is that 'the present queen of France' is *directed* at the present queen of France. That there is no object satisfying the description doesn't threaten its 'directedness' in this sense. Descriptions that have a denotation also have this kind of 'directedness', though perhaps it is easier to overlook since

there is a denoted object to hand. Crucially, that definite descriptions are directed in this way does not depend on there being some object that satisfies the description; the directedness of a definite description does not depend on (and is in some sense prior to) what it refers to.¹³

The same plausibly holds for what we might call ‘definite beliefs’ (e.g. that the cat in the kitchen has black fur), unlike ‘indefinite beliefs’ (e.g. that a cat has black fur). Moreover, a belief being directed in this way does not depend on there actually being a (unique) cat in the kitchen. The same holds, allowing for the appropriate substitutions, for other attitudes (desires, fears, etc.). We can distinguish between definite and indefinite fears, definite and indefinite desires, and in each case can identify a kind of directedness that these definite attitudes have.

The observation that definite representations are antecedently directed in this way lights a path to an attractive approach to aboutness in the face of the problems of the many. Roughly, the proposal is that this kind of *directedness*, which is not explained by one or more objects being selected from the list, is the key to answering HITTING TARGETS. Ordinary utterances (like ‘the cat on the mat has black fur’) manage to be about what they are about not in virtue of successfully selecting some prior object(s) from the list, but merely by being directed at *the cat on the mat*.

Notice just how radical this proposal is: entries on the list do not, either individually or in collections, feature in the explanation of aboutness. This sets my preferred account apart from some proposals in the literature by García-Carpintero (2010), Keefe (2010), and Weatherson (2003: 489-490). García-Carpintero, Keefe, and Weatherson retain some list-based resources in their accounts, even though they step away from *some* traditional referentialist claims, notably when accounting for the apparent possibility of quantifying into intentional attitude reports. Instead, I propose to eschew entirely appeals to the list when explaining aboutness and answering HITTING TARGETS, though exactly how these issues interact with quantification will be discussed in detail below.

It is also worth noting that one might simply insist that the directedness in play here is intrinsically and irreducibly intentional and that it is a mistake to attempt to reduce the directedness to something non-intentional.¹⁴ But for those of us who still hope for a more informative explanation of aboutness and for a satisfying answer to HITTING TARGETS, I will discuss what I take to be the most promising avenue for developing a theory of this kind of aboutness.

When developing a listless account of aboutness, it is natural to look to theories developed with an eye for handling cases in which the relevant representations

¹³And thus, Russell (1905: 487): ‘There is no backward road from denotations to meaning’. Note that when Russell says ‘meaning’ here, he does not mean ‘meaning’ in the sense of Fregean senses. See the two paragraphs following this remark in Russell (1905: 487-488).

¹⁴Thanks to an anonymous referee for making this option salient.

apparently fail to refer. Unsurprisingly, accounts of so called 'empty' representations (representations that appear to be about things that do not exist) often seek to explain directedness in a way that does not require appealing to entries on the list, since it seems on the face of it that the list should exclude such things as witches, the fountain of youth, and so on. My approach to aboutness recent discussion (which I will take as a starting point here) certainly grew out of a focus on empty cases.

5.3 The Co-aboutness First Framework

My (Sandgren 2019; 2021) recent work on intentionality provides us with a foundation for building a listless account of aboutness with an eye on the intentional problem of the many.

One important feature of the proposal is that the aboutness of representations is not explained *in terms of* their intentional objects; singular representations have a directness that is not explained in terms of, or grounded in, a relation to an intentional object. It is not a target-first account. In this respect, the approach resembles some other accounts of singular aboutness. For example, Jeshion's (2002: 67) account that an attitude's or thought's being directed at a particular object is a matter of having a characteristic cognitive role for the representer. It also resembles so called 'phenomenal intentionality' theories like those defended by Farkas (2018) and Kriegal (2008) according to which the aboutness of a representation depends on its phenomenal features rather than its relation to some particular objects. Like the co-aboutness first proposal, these accounts are adapted to apply to cases of empty representation because, roughly speaking, thinking about a particular witch does not require a relation to an actual flesh and blood witch on the list.

One thing that sets Sandgren's approach apart is the contention that the aboutness of a given representation is determined by relationships *between representations*. In particular, the idea is that what a representation is about is determined by which representations it is co-intentional with, that is, which representations it is about the same thing as. On this story, representations get their intentionality in groups. So instead of a set of representations being bound together because they all have the same kind of intentionality, they have the intentionality they have because of their being about the same thing as a set of representations. Their intentionality depends on the broader representational context.

This means that the resulting account of aboutness relies on an account of the co-aboutness relation (the relation of being about the same thing) that stands between representations. If the resulting account is going to eschew list-based explanatory resources, the account of co-aboutness must itself not appeal to relations to entries on the list. I have proposed my own account of the co-aboutness relation based on metarepresentational beliefs that fits that bill (Sandgren 2019), but there are some

other options for which account of co-intentionality to plug in here.¹⁵

The intentional objects (the things representations are about) as I characterise them (Sandgren 2021: n.9), are specified simply in terms of sets of co-intentional representations. The thing that a set of representations is about just is the thing that this or that set of representations are all about, there is nothing more to them. The 'intentional objects', such as they are, depend on the sets of representations and their aboutness; the intentional objects are nothing but the thing that a set of representations is about. These intentional objects are simply supposed to play the role of being the common subject matter tracked across representations that are about the same thing. Crucially for what follows, they are not to be identified with truth makers or with independently individuated entries on the list.¹⁶

5.4 Aboutness and the world

This has been a brief outline of overarching framework. But this leaves many important questions unanswered. I characterise a distinctive set of intentional objects which are tied to sets of representations. But it is left largely unclear how they relate to objects in the world. Even if intentional objects may depend, in some sense, on sets of representations and their intentional features, trains, cats, trees, and other objects 'in the world' do not. What is more, it is plausibly the object in the world that makes truth-apt representations true. Relatedly, when we make true claims like 'there is only one capital of Finland', it seems like they are made true by objects in the world.

Setting aside the problem of the many for a moment, target first accounts of aboutness have a natural story to tell about the relationship between aboutness, quantification, counting, and truth; independently individuated entries on the list are characterised as truth makers, the things over which we quantify, the things we count, and intentional targets. By decoupling aboutness from the list, I cannot so cleanly connect aboutness, truth, quantification, and counting. Things are going to have to get messy. My answer to HITTING TARGETS does not depend on the list and is not threatened by the world's generosity. But how can we harmonise this approach to familiar notions of truth, quantification, and counting, especially when answering questions like COUNTING CATS and ACHIEVING TRUTH.

¹⁵See, for instance, the accounts of co-intentionality from Donnellan (1974), Sainsbury (2010), Crane (2013: 165), or Pagin (2014).

¹⁶It is worth noting that on this proposal, there may still be rough or indeterminate boundaries of intentional targets. It may be indeterminate whether or not a given pair of representations are co-intentional and therefore the boundaries of intentional objects will be somewhat indeterminate. I think this is a virtue rather than a vice of the proposal. It promises to resolve some of the thornier issues in the vicinity of the problem of the many without surrendering the intuitive thought that what we talk and think about is often fuzzy and indeterminate.

1. How does quantification over intentional targets, especially in intentional contexts, work on the listless approach to aboutness?
2. On the listless proposal, how does aboutness (and the corresponding answer to HITTING TARGETS) relate to truth conditions (and ACHIEVING TRUTH)?
3. How does the listless approach handle counting questions like COUNTING CATS?

In what follows I directly address these questions. Through the discussion that follows I will make it clear exactly how one's approach to answering HITTING TARGETS can be separated from the tools used to answer ACHIEVING TRUTH and COUNTING CATS.

5.5 'The Things in the World'

At this stage it will be illustrative to consider a natural objection to my approach to aboutness. Plausibly, we are able to think about things *out there in the world*, but I have resigned myself to using only the semantic resources compatible with failures of reference – in fact, I explicitly motivated the view *via* denoting phrases that have no denotation! But how could a notion of aboutness (as mere directedness) equip 'aboutness' to play its intended theoretical role of connecting us to the things in the world? If the list doesn't figure in explaining how our utterances manage to be about particular things, how could ordinary thought and talk be said to 'connect up' with the world and the things in it in any meaningful sense?

A defender of the my framework should respond as follows. Recall that 'The world' and 'things in the world' are just more representations. So when addressing the question of whether we can think about 'things in the world' we need simply check to see if we can produce or indicate representations that are about the same thing as 'the things in the world'. If 'the things in the world' is about any object or set of objects at all, then we can think and talk about them; indeed, we just did by producing 'the things in the world'. What is happening here is a subtle reinterpretation of the question 'can we talk and think about things in the world?'. We can reply in the affirmative. In fact, the fact that the question can be sensibly asked (and the question is about the things in the world) is more or less sufficient to make the correct answer 'yes'. We have a test for what it takes for a representation to be about the things in the world. What it takes is for that representation to be about the same things as 'the things in the world'.

5.6 Quantification over Intentional objects

But perhaps what the questioner intended was for the question 'can we talk and think about things in the world?' to be interpreted in a different way. Perhaps the

questioner really wanted to know, for some object in the world, can we talk about *it*.

One might think that when there is a wide-scope quantification over objects in the world, and we then quantify into an intentional context, the intentional object is, in a sense, forced to be an object in the world rather than one of my intentional objects. For instance, 'there is a cat that Jill has beliefs about' or 'Jill has beliefs about every cat in the house'. Those who understand aboutness referentially can identify the intentional targets of our thought and talk with their referents, and so can understand quantification straightforwardly in terms of quantification over referents selected from the list. But those who defend listless aboutness will need a different story. Insofar as the current proposal involves eschewing 'the list', how can we make sense of quantification over the targets of our thought and talk?

This is a serious concern. Indeed, García-Carpintero (2010: 358) and Weatherson (2003) suggest that a major cost of eschewing list-based resources entirely is that it compromises our ability to account for quantification over the targets of our thought and talk.

The core of my response is the observation that 'quantifying over' something is just another way of talking about something. We do not reach out and select *things in themselves* by quantification, however wide the scope. We quantify over the very same things we talk about — cats, trees, clouds, and so on. In line with this, once we have established the aboutness facts, we get intentional targets that we can quantify over *for free*. These objects are not supplied by the list but are a kind of shadow of the aboutness facts, so the aboutness facts determine the facts about the individuation of intentional objects. This is a kind of 'easy ontology' approach to intentional objects and if this kind of suggestion is right, it lights a path to a satisfactory answer to questions concerning quantification over intentional objects. It is these intentional objects, rather than entries on the list, that we quantify over when we quantify into intentional contexts.

Of course this proposal will need to be supplemented with an account of how we can count intentional objects, when assessing claims like 'Jane is talking about three cats'. I will leave discussion of these counting questions to section 4.8 below.

5.7 Aboutness and Truth

Next, how does the listless conception of aboutness interact with truth conditions and truth values? Consider 'the cat in the kitchen is black' and suppose it is true. According to the proposed approach to aboutness, this is about the cat in the kitchen but that fact does not depend on there being a cat in the kitchen (let alone its being black). But that claim's being true *does* depend on the list; for it to be true, there really has to be a black cat in the kitchen.

Although I am proposing that entries on the list do not figure in our explanation of aboutness, we need not (and should not) eliminate them from our account of truth values. One's account of truth values should appeal to whether there are things that are as they're represented as being. Recall the distinction between truth conditions and aboutness I appeal to in section 5.1 and footnote 3 above. The aboutness of a representation need not be taken to supervene on its truth conditions (let alone truth values). This, I think, is a crucial part of the proposed account. Those who adopt the proposed approach should distinguish the object or objects that make a representation true and what that representation is about. The object that makes 'the cat in the kitchen is black' true (which *really is black*) will be distinct from what that sentence is about, the shadow of the aboutness facts. We should then understand the truth conditions of the claim in a way that does not pertain to what it is about. In this case the truth conditions are relatively simple: 'the cat in the kitchen is black' is true just in case there is just one cat in the kitchen that is black. Those truth conditions do not require any one object to be the truth maker across all the situations in which the truth conditions are satisfied.

But aren't sentences and beliefs involving names different? What about rigid designators and beliefs that have a similar 'direct reference' flavour? Again there are a number of what we might call 'generalist' approaches to truth conditions in the offing (Lewis 1970; 1986; 1981; 2009: 33) mostly developed in response to so called 'Frege cases'. The idea is to tie truth conditions of representations (including those involving names and demonstratives) to the satisfaction of a *role* rather than in terms of the actual *realiser* of the role. For instance, Pierre's belief that London is pretty is true just in case there is a Londonizer (something that plays the role Pierre associates with London) that is pretty. This generalist approach to truth conditions is controversial, though it has some appealing features and I have defended it elsewhere (Sandgren forthcoming). In any case, the defender of my account of aboutness has a suitable account of truth conditions ready to hand.

Another option is to deploy a supervenience account of the relevant truth conditions here, perhaps drawing on McGee and McLaughlin (2000). The idea would be that, when evaluating claims like 'the cat in the kitchen is black', the claim is true just in case all the members of the CAR are black. The idea would be that this is an instance where it is really super-truth that is at issue. Of course, this requires that the supervenience answer to ACHIEVING TRUTH is sufficient, but in either case this sort of approach to truth conditions does not hinder the proposed approach to aboutness. This also reveals one reason why a relatively thorough discussion of the method of supervenience was called for in this article. The discussion was required to show that supervenience tools do not suffice for an account of aboutness and that, therefore, the account defended here is not a rival to supervenience.

Things are more complicated when we bring quantification over intentional objects into the mix. Consider, for example, ‘the cat that Jane is talking about is black’ and suppose that it is true. How, without appealing to the list, are we to make sense of the truth conditions of utterances like that? It seems like the intentional object of Jane’s belief must form part of the explanation for the truth of that claim. To have black fur an object needs to *really* be a certain way. Similarly, to be thought about by Jane is one thing, while simply being represented to be thought about by Jane is another.

When we quantify over the intentional targets, we are again not talking about the truth makers of that sentence. When we evaluate that claim, we have thoughts about *the intentional target of Jane’s talk* and our representation is directed at that intentional object. Of course, to make that claim true, an object needs to really be what Jane is thinking about, but that object will not be the intentional object of ‘the object Jane is thinking about’ as we represent it (recall the distinction between intentional objects and truth makers). Again, simply ascribing relations of aboutness to an object and a representation does not let us think about things in themselves, rather we simply arrange another layer of intentional objects. This means that it is benign that the intentional targets need not have the features that entries on the list have and which would allow them to play the dual role of intentional targets *and* truth makers. This implies that there will not be an elegant mapping between truth conditions (or truth makers) and aboutness.

This shiftiness is inelegant, since the truth maker is, in a sense, intentionally elusive. We can never pick it out without falling back onto another representation.¹⁷ I am happy to accept that the semantic theories linking ordinary thought and talk (including ordinary quantification) to truth makers are likely to be somewhat inelegant.¹⁸ Any adequate account of quantification over intentional objects and how intentional targets relate to truth conditions and truth makers will be messy, but this messiness might be necessary, or at least involve a cost that proponents of the view on aboutness I have been discussing should be willing to pay.

5.8 Counting

How does the listless approach to aboutness tie into counting questions? Ultimately, what is required is an answer to COUNTING CATS that harmonises with the proposed approach to the intentional problem of the many.

Consider ‘there is exactly one cat on the mat’. This claim seems true when, as we would ordinarily say, there is one cat on the mat. As discussed above, deploying supervenience tools can help here and McGee and McLaughlin (2000)

¹⁷Here I am influenced by Lewis’ ‘Ramseyan Humility’ (2009).

¹⁸See Manley (ms) on closely related issues, especially section 8.

take exactly this line. The idea is that for each of member of the CAR, if that object is a cat on the mat, none of the other members of the CAR is a cat on the mat. The idea is that this would make 'there is exactly one cat on the mat' super-true.

Another option is to claim that we often count by relations other than strict identity. Perhaps the best known version of this kind of approach is due to Lewis (1993), though an earlier example of a view of this kind is due to Geach (1967a). According to Lewis, we can appeal to relations of *almost* identity which are based on massive overlap. The many cats on the mats are distinct but *almost* identical since they massively overlap. When it comes to cats, the differences between the many are negligible. So, on one good reading, 'there is exactly one cat on the mat' is true just in case there is a cat on the mat and every cat on the mat is almost identical to all the cats on the mat.¹⁹ This is what I mean when I say 'counting by' relations other than identity; it means that the uniqueness condition (or number condition) is formulated in terms of some relation other than strict identity. On this reading, 'there is exactly one cat on the mat' is true when there is, as we would ordinarily say, one cat on the mat since all the cats stand in some other relation to each other.

Another way to implement the idea of counting by relations other than identity is due to Geach. Geach claims that we count by relations of partial indiscernibility and that individuation and counting are relative to a sortal. On Geach's proposal,

“Same cat” is a relation of partial indiscernibility, restricted to respects of comparison somehow associated with the term 'cat', and discernibility by just a few hairs doesn't count. “Same lump of feline tissue” is a different relation of partial indiscernibility, and a more discerning one. (Lewis 1993: 175)

The relevant sortals embed semi-stable information about which differences matter when answering questions about individuation; we count sheep in one way and solar systems in another. These relations of being the same F as, where 'F' is a sortal, is another kind of relation other than identity by which we could count.

In any case, these moves are independent of one's approach to the *intentional* problem of the many. Recall that, none of the members of the CAR are the intentional objects and neither are the many cats that stand in the crucial relation. In this sense, the proposed account of aboutness is compatible with a number of answers to COUNTING CATS. For what it is worth, I think that accounts that involve counting by relations other than identity are the most attractive here, though I think that

¹⁹See López de Sa (2014) for an especially helpful discussion of Lewis on counting questions. A reason for pause here – evident in this discussion and throughout the literature – is that COUNTING CATS might be better understood as a question about the *metaphysics* of cats, and so not best answered by rescuing the truth of sentences like 'there is exactly one cat on the mat'. See, for example, the discussion in McGee and McLaughlin (2000: 140-142).

Lewis' view of this kind of relation is crucially limited. For more on expanding our understanding of the relation other than identity by which we count, see López de Sa (2014).

There is one more sense in which our account of aboutness should be informed by the problem of the many and particularly the work on counting by relations other than identity. According to my proposal, intentional objects are tied to sets of representations so there are as many intentional objects as there are sets of representations that are about the same thing. But this story needs to be complicated somewhat in any case. Say I hold two copies of the same newspaper and ask 'same or different?'. Before you know how to answer, you need to know what kind of thing I am asking about. They are the same issue (I am only holding up one issue) but they are different copies (there are two physical objects, one in each hand). Suppose we ask, pointing to the thing in my left hand, 'which representations are about that?'. Different answers are appropriate depending on what *kind* of thing the representation of the copy in my left hand is about. A representation could be directed at the copy or it could be directed at the issue, and depending on which, it will be co-intentional with different representations. So, in a sense, relations of co-aboutness will need to be made sensitive to this kind of distinction. In some way or another, the *kind* of thing at which a representation is directed needs to be accounted for in any co-aboutness based account of aboutness. I briefly discuss on how one might augment my account of co-aboutness (Sandgren 2019) so as to make it sensitive to the *kind* of thing being thought or talked about in my dissertation (Sandgren 2016, 78-79).

The proposal I prefer here mirrors the 'counting by relations other than identity' idea from Geach. Relations of co-aboutness (and therefore intentional objects) should be made relative to some sortal, which encodes, broadly speaking, the kind of thing that the intentional target is. On this story, two representations could be about the same *issue* without being about the same *copy*. The copy and the issue qua intentional objects are distinct, with distinct corresponding sets of co-intentional representations.

5.9 Concluding Remarks

List-based approaches to representation have been dominant in the literature since the 1980s, largely due to the influential work of Donnellan (1970), Putnam (1975), Kripke (1980) and others. I do not expect to sway many dyed-in-the-wool supporters of list-based approaches, though perhaps I have said enough to tempt those readers who have not yet firmly made up their mind on the matter.

Acknowledgements

This paper began life coauthored with Maegan Fairchild and she played a particularly central role in the first part of the paper (sections 1-3). I want to express my gratitude for her brilliant input and ridiculously fun philosophy chats on the material in this paper. I would also like to thank Pär Sundström, Martin Abreu Zavaleta, Bram Vaassen, Luise Mirow, Caroline Torpe Touborg, Per Algander, Torfinn Huvenes, Daniel Nolan, and two anonymous reviewers for *Ergo* for helpful feedback on earlier versions.

References

- Bernstein, S. (2021). Could a middle level be the most fundamental? *Philosophical Studies* 178(4), 1065–1078.
- Brown, J. D. K. (2016). Natural objects. *Journal of the American Philosophical Association* 2(2), 254–271.
- Carnap, R. (1950). Empiricism, semantics, and ontology. *Revue Internationale de Philosophie* 4(11), 20–40.
- Chalmers, D. (2002). The components of content. In D. J. Chalmers (Ed.), *Philosophy of Mind: Classical and Contemporary Readings*. Oxford University Press.
- Crane, T. (2013). *The Objects of Thought*. Oxford University Press.
- Donnellan, K. S. (1970). Proper names and identifying descriptions. *Synthese* 21(3-4), 335–358.
- Donnellan, K. S. (1974). Speaking of nothing. *Philosophical Review* 83(1), 3–31.
- Fairchild, M. (2019). The barest flutter of the smallest leaf: Understanding material plenitude. *Philosophical Review* 128(2), 143–178.
- Fairchild, M. (2020). Varieties of plenitude. *Philosophy Compass* 15(3), 1–11.
- Farkas, K. (2008). Phenomenal intentionality without compromise. *The Monist* 91(2), 273–93.
- Fine, K. (2014). Truth-maker semantics for intuitionistic logic. *Journal of Philosophical Logic* 43(2-3), 549–577.
- Fine, K. (2020). Yablo on subject-matter. *Philosophical Studies* 177(1), 129–171.

- García-Carpintero, M. (2010). Supervaluationism and the report of vague contents. In R. Dietz and S. Moruzzi (Eds.), *Cuts and Clouds: Vagueness, its Nature and its Logic*. Oxford University Press.
- García-Carpintero, M. (2020). Co-identification and fictional names. *Philosophy and Phenomenological Research* 101(1), 3–34.
- Geach, P. T. (1967a). Identity. *Review of Metaphysics* 21(1), 3–12.
- Geach, P. T. (1967b). Intentional identity. *Journal of Philosophy* 64(20), 627–632.
- Hawthorne, J. (2006). *Metaphysical Essays*. Oxford University Press.
- Hawthorne, J. and A. McGonigal (2008). The many minds account of vagueness. *Philosophical Studies* 138(3), 435–440.
- Keefe, R. (2010). Supervaluationism, indirect speech reports, and demonstratives. In R. Dietz and S. Moruzzi (Eds.), *Cuts and Clouds: Vagueness, its Nature and its Logic*. Oxford University Press.
- Korman, D. Z. (2015). *Objects: Nothing Out of the Ordinary*. Oxford University Press UK.
- Kriegel, U. (2008). The dispensability of (merely) intentional objects. *Philosophical Studies* 141(1), 79–95.
- Kripke, S. A. (1980). *Naming and Necessity*. Harvard University Press.
- Lewis, D. (1970). General semantics. *Synthese* 22(1-2), 18–67.
- Lewis, D. (1981). What puzzling pierre does not believe. *Australasian Journal of Philosophy* 59(3), 283–289.
- Lewis, D. (1983). New work for a theory of universals. *Australasian Journal of Philosophy* 61(4), 343–377.
- Lewis, D. (1984). Putnam's paradox. *Australasian Journal of Philosophy* 62(3), 221–236.
- Lewis, D. (1986). *On the Plurality of Worlds*. Wiley-Blackwell.
- Lewis, D. (1993). Many, but almost one. In K. Cambell, J. Bacon, and L. Reinhardt (Eds.), *Ontology, Causality and Mind: Essays on the Philosophy of D. M. Armstrong*, pp. 23–38. Cambridge University Press.
- Lewis, D. (2009). Ramseyan humility. In D. Braddon-Mitchell and R. Nola (Eds.), *Conceptual Analysis and Philosophical Naturalism*, pp. 203–222. MIT Press.

- López de Sa, D. (2014). Lewis vs Lewis on the problem of the many. *Synthese* 191(6), 1105–1117.
- Lowe, E. J. (1982). The paradox of the 1,001 cats. *Analysis* 42(1), 27–30.
- Lowe, E. J. (1995). The problem of the many and the vagueness of constitution. *Analysis* 55(3), 179–182.
- Manley, D. Keeping up appearances: A reducer’s guide.
- McGee, V. and B. P. McLaughlin (2000). The lessons of the many. *Philosophical Topics* 28(n/a), 129–152.
- McGrath, M. (2005). No objects, no problem? *Australasian Journal of Philosophy* 83(4), 457–486.
- Merlo, G. (2017). Multiple reference and vague objects. *Synthese* 194(7), 2645–2666.
- Openshaw, J. (2020). Thinking about many. *Synthese*, 1–20.
- Pagin, P. (2014). Intersubjective intentional identity. In M. García-Carpintero and M. Genoveva (Eds.), *Empty Representations: Reference and Non-Existence*, pp. 91–113. Oxford University Press.
- Putnam, H. (1975). The meaning of ‘meaning’. *Minnesota Studies in the Philosophy of Science* 7, 131–193.
- Rettler, B. (forthcoming). Mereological nihilism and puzzles about material objects. *Pacific Philosophical Quarterly*.
- Russell, B. (1905). On denoting. *Mind* 14(56), 479–493.
- Sainsbury, R. M. (2010). Intentionality without exotica. In R. Jeshion (Ed.), *New Essays on Singular Thought*.
- Sandgren, A. (2016). *Cruel Intentions: An Essay on Intentional Identity and Intentional Attitudes*. Ph. D. thesis, The Australian National University.
- Sandgren, A. (2019a). A metarepresentational theory of intentional identity. *Synthese* 196(9), 3677–3695.
- Sandgren, A. (2019b). Puzzling pierre and intentional identity. *Erkenntnis* 84(4), 861–875.
- Sandgren, A. (2021). Turning aboutness about. *Pacific Philosophical Quarterly*.
- Sandgren, A. (forthcoming). Private investigators and public speakers. *Journal of the American Philosophical Association*.

- Sider, T. (2009). Ontological realism. In D. J. Chalmers, D. Manley, and R. Wasserman (Eds.), *Metametaphysics: New Essays on the Foundations of Ontology*, pp. 384–423. Oxford University Press.
- Sider, T. (2011). *Writing the Book of the World*. Oxford University Press.
- Toner, P. (2008). Emergent substance. *Philosophical Studies* 141(3), 281–297.
- Unger, P. (1980). The problem of the many. *Midwest Studies in Philosophy* 5(1), 411–468.
- Weatherson, B. (2003). Many many problems. *Philosophical Quarterly* 53(213), 481–501.
- Williams, J. R. G. (2007). Eligibility and inscrutability. *Philosophical Review* 116(3), 361–399.
- Williams, J. R. G. (2015). Lewis on reference and eligibility. In B. Loewer and J. Schaffer (Eds.), *A companion to David Lewis*, pp. 367–382. Wiley-Blackwell.
- Williams, J. R. G. (2020). *The Metaphysics of Representation*. Oxford University Press.
- Yablo, S. (2014). *Aboutness*. Princeton University Press.