

Generality to Generalised

Introduction


Gareth Evans (1982) and many following him have argued that it is a constraint on a good theory of thought that it make sense of the following fact: If a subject can think that a is F and can think that b is G then the subject can think that a is G and can think that b is F. Not all representational systems are general in this way and contrasting cases are illuminating. “When we attribute to the brain computations whereby it localises the sounds we hear, we *ipso facto* ascribe to it representations of the speed of sound and of the distance between the ears, without any commitment to the idea that it should be able to represent the speed of light or the distance between anything else” (Evans 1982, fn. 22, p 104). Here is another example. Look at squares labeled ‘X’ and ‘Y’ in fig. 1 – a Farnsworth Munsell Hue Test¹ array, which is an array of colour patches, many of which are very similar to each other, some of which match each other, and many of which are clearly different in colour. To my eye, they look the same in colour – I *visually represent* them as the same. When I cast my eyes to the square labelled ‘Z’, I’m pretty sure it is the same too, but I’m not visually presented with sameness of colour. Rather, I draw on memory and quick flits back and forth from X and Z. I can’t get a good, direct look at them at the same time. I *can* get them both into my periphery, but when I do, I no longer find myself in a good position to tell if Z is the same colour as X. So, just because I can visually represent each of X and Y and Z and I can visually represent that X is the same colour as Y, I am not thereby able to visually represent that Z is the same colour as X.²

In light of considerations like those above, Evans concluded that thoughts are ‘general’ but perceptual states are not. On this basis, he argued that thoughts must have a

¹ <https://www.xrite.com/hue-test>

² Davies (1992) offers another example of a failure of generality: a creature might have a hawk-diving-on-beetle detector, but no beetle-diving-on-hawk detector. Beck (2012) argues that ‘analogue magnitude states’ also violate generality. The example of the colour patches is inspired by Heck (2007).

Visually, X is represented as the same colour as Y



Visually, is X is represented as the same colour as Z?

kind of structure not found in perceptual states, and for Evans, that difference in structure is to be located in the *contents* of the states in question. This was one reason that Evans took thoughts to have *conceptual content* and perceptual states, which do not exhibit generality, to have *non-conceptual content*. But this isn't the only way to go. Since Evans's writing, an important distinction has been drawn. Although many philosophers of mind are sympathetic to the conceptual/non-conceptual distinction, there is disagreement over the nature of the non-conceptual. *Content* non-conceptualists follow Evans in holding that being non-conceptual is a property of the *contents* of certain representational states. *State* non-conceptualists hold that being non-conceptual is a property of some non-content aspect of the representational *states* themselves. For example, one might, following Fodor, argue that thinking involves a linguistically structured Language of Thought but vision doesn't.³

³ One suggestion in the literature is that the mental representations in cases of belief are sentence-like and the representations in cases of perception are picture- or map-like. See Camp (2007) for a helpful discussion of map-like mental representations. See Fodor (2008), especially chapter 6 for a treatment of perceptual representation in terms of 'icons'. My own view (Grzankowski 2015) is that pictures have propositional content but differ from sentences at the level of the vehicle of content.

I agree with Evans that thoughts must have a kind of structure that makes sense of the generality of thought but I don't find either of the above ways of telling the story general enough. As I'll argue below, following Evans leads to an overly complicated story about content and following Fodor leads to an overly specific and non-law-like story (something Fodor himself seems to agree with). I'll present a simple alternative story that draws on recent work on non-propositional attitudes and on recent work on the unity of the proposition. The result is an explanation of generality at the level of folk-psychological states rather than at the level of content or vehicle.

A few preliminaries are in order. Evans introduced generality as a constraint on a good theory of thought, a constraint that is met by mature, human thinkers. According to the constraint, if S can think that a is F, then, for every α and Φ for which $[\alpha \text{ is } F]$ and $[a \text{ is } \phi]$ are thinkable, if S can think about α , S can think $[\alpha \text{ is } F]$, and if S can think about Φ , S can think $[a \text{ is } \phi]$.⁴ For example, if a subject can think the thoughts that Sally is tall and that Mary is short, then she can think the thoughts that Sally is short and that Mary is tall: generality, first and foremost, says that the having of certain abilities entails the having of other abilities. The Generality Constraint has been used to establish conclusions about the architecture of mind, the nature of concepts, and the nature of contents. We are presently interested in the metaphysics of thought and what it must be like to meet the Constraint.

Generality concerns idealised thought and thinking. It may be that one who can think the thoughts that a is F and that b is G cannot form the thought that b is F because he has, for example, sustained a blow to the head or because the thought would be too psychologically traumatic. But, typically (as Heck puts it)⁵ or ideally (as Evans puts it),⁶

⁴ Dickie (2010, p. 501) defends putting the constraint this way. For near variants see Evans (1982, p.104) and Peacocke (1992, ch. 2).

⁵ 2007, p. 121-122.

⁶ 1982, p. 105.

one who can think that a is F and that b is G can think that b is F.⁷ In contrast, even when all is going well, the examples above of the distance between the ears and the squares X, Y, and Z still flout the Generality Constraint.

Let me also flag that generality is not the same as the so-called ‘productivity’ of thought. According to the productivity thesis, there are an infinite number of thoughts that thinkers are in principle capable of entertaining. A number of philosophers have concerned themselves with explaining how we could have an unbounded capability given finite means. That issue is distinct from generality though perhaps explained in the same way.

Finally, I am not presently concerned with whether or not there should be categorial restrictions placed on thought. One might suggest that one cannot think the thought that seven is tall or the thought that Churchill is prime for there are no such thoughts to be had. Such a position is an artefact of the way, for example, Peacocke (1992) thinks of thought. But I see only theory-driven reasons for adopting such a position. We want a smooth account of what goes on in actual thinkers, and I see no reason to place a bar on the thought that, say, Churchill is prime.⁸ But since these are controversial issues, focus will be on thinkable thoughts, whichever those might be.

The Standard Theory

We are working under the assumption that being general is in thought’s very nature. Accordingly, theorists such as Evans (1982), Peacocke (1992), and Heck (2007) have demanded that a theory of thought meet the Generality Constraint *in a constitutive way*. That is, they demand an account of thought that answers the question ‘What is it about thought such that a subject who can think the thoughts that a is F and that b is G can thereby think the thoughts that a is G and that b is F?’.

⁷ See also Fodor (1998) who discusses generality under the label ‘specificity’.

⁸ See Camp (2004) for further discussion and a defense of the view that one should not rule out such cases as instances of thoughts.

Presently, I want to offer the core of what I'll call 'the standard theory' of thought. The kind of view was first suggested by Evans (1982),⁹ and Peacocke (1992) offers a related view in terms of concepts and their relation to the contents of complete thoughts that builds out Evans's suggestion.¹⁰ Their developments of the position come with various commitments that would be a distraction at present, so I'll try to distill the crucial commonalities.

Evans notes in his original discussion that it's hard to see how the Generality Constraint could be met if thoughts weren't structured, and it's tempting to think that the needed structure is to be found, if anywhere, in the *contents* of thoughts.¹¹ On a Fregean model, contents comprise concepts that must be exercised if one is to be in a state with that content. The content itself is a structure of concepts. To think the thought that a is F is to stand in a certain kind of relation to the Fregean content $\langle \mathbf{a}, \mathbf{F} \rangle$ which in turn requires exercising the constituent concepts \mathbf{a} and \mathbf{F} and combining them. Suppose, for example, that one believes that a is F. On the Fregean view presently under consideration, to believe that a is F is to stand in a two-place relation, the relation designated by 'believes', to the proposition that a is F. But to stand in that relation to that proposition, one must exercise the concept \mathbf{a} and must exercise the concept \mathbf{F} . And one could tell a similar story in a Russellian vein. Let a proposition be a structure not of concepts but of objects, properties,

⁹ See especially Ch. 4 sections 3 and 4. In an appendix, Peacocke (1992) offers a comparison between his view and Evans's, arguing (persuasively I think) that they are by and large in agreement.

¹⁰ See also Heck (2007) for further discussion, especially in the context of the non-conceptual.

¹¹ Standard theorists typically take thoughts to be two-place relations between thinkers and their contents. It is indeed hard to see, on such a view, where to locate the structure required to meet the Generality Constraint if not in the contents of thoughts. Peacock takes it as an assumption that 'attitudes are relations to complex contents, composed in a distinctive way from concepts possessed by the thinker' (1992, p. 44). In discussing psychological explanation and Frege's puzzle, Heck (2012) tells us the following:

Since Fred's belief that Mark Twain has died and his belief that Sam Clemens has died are different beliefs, they must have different contents, whence the contents must be individuated more finely than Russellian propositions are. This move is an extremely natural one, often tacitly made. **What lies behind it, it seems to me, is the thought that beliefs just are relations between thinkers and contents.** And what lies behind that thought, I suggest, is the view that psychological explanation (in the sense in which we are concerned with it) is intentional explanation, that is, explanation in terms of the contents of psychological states as opposed, say, to the neurological properties of such states. (p. 145, emphasis added)

More below on intentional explanation.

and relations. Now add that bearing the relation designated by 'believes' to it requires representing (or perhaps being acquainted with) the constituents of the content.

Notice that on the views being sketched, there is a two-part story. Thoughts are structured by virtue of having structured contents – they are two-place relations between thinkers and propositions. But by itself this won't meet the Generality Constraint. Contents have constituents, but they are also a 'unity' – they are something over and above their constituents. How is being related to a unified content that has constituents supposed to help us meet the Generality Constraint? The proponent of structured content must also have some operation that "sees past" the unity, as it were. If our theory of thought gives us no more than two-place relations to propositions, the fact that those propositional wholes have constituents is useless with respect to generality: a subject who stands in R to $\langle \mathbf{a}, \mathbf{F} \rangle$ and to $\langle \mathbf{b}, \mathbf{G} \rangle$ may have no way of getting onto $\langle \mathbf{b}, \mathbf{F} \rangle$. At a level of abstraction that suits our present needs, we should add that being able to have thoughts with structured contents requires the having of a series of abilities (and indeed this is exactly what one finds in Evans: he tells us that he 'should prefer to explain the sense in which thoughts are structured, not in terms of being composed of several distinct *elements*, but in terms of their being a complex of the exercise of several distinct conceptual *abilities*' (1982, p. 101)). One must have some ability to "get onto" a and "onto" F and one must have a mechanism or ability for combining the constituents in appropriate ways – a 'mode of combination' as Peacock calls it (p. 24). Peacocke takes it to be an implicit feature of possessing a predicate concept such as F that a 'thinker knows what it is for an *arbitrary* object to fall under F' (p. 43, my emphasis).¹² ¹³ But one needn't build this into concept possession. It would suffice

¹² See also pages 23-25 where Peacocke introduces and more fully explicates his views of concepts and concept possession.

¹³ When Peacock says 'arbitrary object' he means to restrict the class to those objects, concepts of which, fall under the predicate concept's 'range of significance' (p. 42). That restricted set of objects is such that concepts of them can be combined with predicate concepts to form representations that have truth-conditions. Whether a theory of thought should make such restrictions seems to me an open question, but one that needn't be taken up here. As I said above, the present discussion will be limited to thinkable thoughts, whichever those might be.

to hold that one must have a general ability to combine concepts. But however one wishes to spell out those details, according to the view under consideration, a subject who can think the thought that a is F and can think the thought that b is G has all that is needed to form the thoughts that b is F and a is G, for the subject is able to exercise all of the constituent concepts (or in the case of the Russellian view, represent the constituent objects and properties) and has the general ability to appropriately combine those constituents.¹⁴

In some ways, I simply agree with the Standard theorists, at least at the level of abilities underwriting our relations to contents. But I struggle to see why they draw further conclusions about the nature of the *contents* themselves. In effect, the structured view of contents (at least in the present debate) *encodes* demands on thinkers. That is, encodes a demand on being related to the contents. This isn't an absurd idea, but it does look to be optional. Focus on the Fregean view. We are asked to hold that contents have concepts as constituents and to then go onto say that those concepts must be grasped or exercised by a thinker who stands in the relation designated by, say, 'believes' to such a content. Contents that lack structure or lack conceptual constituents place no such demands on being related to them. On such a picture, the Fregean contents have built right in that bearing conceptual, propositional attitude relations to them requires the series of abilities that explains generality. And as noted above, this is exactly the sort of picture that then leads to a distinction between kinds of contents – conceptual and non-conceptual. If we take this path, we wouldn't want to attribute to audition or vision such contents since doing so

¹⁴ It is worth flagging the sense in which generality is really 'explained'. Notice that the theory on offer posits a *general* mechanism in light of the Generality *Constraint*. Generality as such is not *explained* in the sense of reduced, analysed, or eliminated on such a theory, nor is such an explanation the target (at least not for Evans and those following him). Rather, the Generality Constraint (it is a *constraint* after all) serves to guide a theory of thought and demands certain structural features that *accommodate* a data point. More on the Language of Thought below, but a similar line of thought appears there: the grammar of the LOT is typically taken to be general in such a way as to accommodate features of thought such as productivity and systematicity.

would then entail demands that audition and vision do not in fact meet. But why take this path?

We ought to resist complicating contents. Once one has the series of abilities – abilities to exercise **a**, **F**, and combine them – why bother building this into *content*?

Part of the point of positing content is to capture representational similarities and relations holding amongst representations such as sentences, perceptions, beliefs, and so on. For example, two speakers can say the same thing by uttering different sentences, perhaps one in English and the other in French. What was said, what the sentences represent, is the *content* of the two distinct sentences. Contents allow us to draw out representational similarities by abstracting away from the specific representational media (in this case, token sentences of distinct languages). Similarly, I might *believe* what someone else *said* – the content of my belief is the same as the content of what was said; they represent the same.¹⁵ This kind of motivation for content extends to perception as well.¹⁶ I might *believe* that things are just as I *perceive* them to be – the contents of my perceptions are contents I believe. Contents are brought in precisely (though perhaps not exclusively) to account for facts like these. If some representations have different *kinds* of content than others with which they interface, familiar and central roles we ask contents to play will at the least be needlessly complicated. If we were to hold that, say, English sentences have a different kind of content than French sentences, a simple account of synonymy and translation in terms of sameness of content would have to be given up. If we were to hold that the contents of thoughts are of a different kind than the contents of

¹⁵ Stalnaker (1990) and Schiffer (2003) both offer helpful discussions. See Grzankowski and Buchanan (2019) for a view according to which propositions simply are those entities that are in common between things that represent things as being the same way.

¹⁶ See Siegel (2011) for a helpful overview. The view that perceptual experiences have content is widely endorsed. For example, Block (1996), Evans (1982), Harman (1990), McGinn (1989), Peacocke (1983), Searle (1983) and Shoemaker (1990) think that perceptual experiences have an intentional component. Theorists such as Byrne (2001), Dretske (1995, 2003), Lycan (1996), McDowell (1994), Pautz (2009, 2010), Siegel (2010), Thau (2002), and Tye (2000) argue for the stronger position that even the phenomenology of perceptual experiences is exhausted by their contents. For general dissent to content views see Alston (1998), Brewer (2011), and Travis (2004). For a view that aims to capture the virtues of both sides of the debate, see Schellenberg (2018).

assertions, then it is hard to see how to make sense of successful communication in terms of content sharing. Similar considerations apply to perception. If we were to hold that the kinds of contents had by perceptions are of a different kind than those had by beliefs, the story about how we can believe what we perceive is more complicated than we might have hoped. Moreover, how what we perceive *justifies* what we believe will also be made more complicated – a theory in terms of the logical relationships holding between contents is no longer straightforward.¹⁷ In short, positing distinct kinds of content upsets major motivations for positing contents in the first place and brings with it a host of issues about how representations with distinct kinds of content are supposed to interface. I can't do justice to these familiar thoughts in the present paper and it is not my position that distinct kinds of content could never be justified or that roles of content just outlined couldn't be recovered. I hope, however, that these brief points help to remind one of a key reason for positing contents and highlight that if we are going to posit distinct *kinds* of contents, specifically for belief and perception, there is a price to be paid. If we can avoid paying it, I think we should. I want to argue that when it comes to the generality of thought, we can have everything we want without multiplying kinds of content. After all, the Standard Theorist seems to have the resources of various abilities such as exercising concepts and a general mode of combination without saying anything about content, so why not skip the problematic step?

Why A Language of Thought Isn't Enough

Maybe there *is* a good reason to complicate contents even if it comes at a cost. This can be brought out by considering an alternative to the Standard Theory.

At this juncture, isn't it obvious that there is an attractive explanation that will not require making any commitments concerning the nature of content? In particular, one

¹⁷ See McDowell (1994) for a theorist who argues from such epistemological considerations to the conclusion that all contents must be conceptual.

might reasonably think that a Language of Thought (LOT) is well positioned to meet our needs.¹⁸ But there are some good reasons for thinking that the Language of Thought isn't up to the task and seeing why helps one see why one might wish to follow Evans, Peacocke, and Heck after all. Perhaps, on balance, complicating kinds of contents is our best bet.

According to the LOT proponent, mental representations – representational vehicles – are sentences in a language. Crucially, like sentences of a public language, LOT sentences decompose into representational parts (terms or words in a LOT) and have a compositional grammar. To think a thought, on this view, is to token a mental sentence that has a semantic value that is determined by the semantic values of the parts of the sentence and their arrangement. To illustrate, suppose that a subject believes that *a* is *F*. According to a LOT proponent, the subject must token a sentence in the head that expresses that *a* is *F*. A sentence that expresses that *a* is *F* does so in virtue of the semantic values of its constituent parts (plausibly, one of which refers to *a* and the other to *F*) and their arrangement. Importantly, for our purposes, the structure of the mental sentence may or may not be mirrored in the content of the state. The view could, for example, mesh structured mental sentences with non-structured contents such as possible worlds propositions. The view is relatively silent about the nature of contents but offers structure elsewhere.

It is important to see that the generality of thought does not simply *fall out* the LOT hypothesis. The theory provides representations with meaningful constituents, but by itself this isn't enough (for reasons we saw above). We also need some operation that guarantees access to the constituents of that structure as well as a general combinatorial operation. With respect to the LOT, it is tempting to think that *compositionality* comes to the rescue here. But notice that compositionality merely tells us that meaningful wholes are composed of meaningful parts put together in an appropriate way. In language, for

¹⁸ I have in mind especially views like that advanced by Fodor (1975) and Field (1978) according to which believing that *a* is *F* requires that one stand in a relation to a sentence in the LOT that expresses that *a* is *F*.

example, a compositional grammar tells us that the meaning of ‘Peter loves Mary’ is determined by the meaning of ‘Peter’, the meaning of ‘loves’, the meaning of ‘Mary’, and the way they are combined. It is a thesis about internal arrangement. Generality, on the other hand, is a requirement on richness: the ability to have certain thoughts guarantees the ability to have others. The LOT has two roadblocks to overcome. First, nothing in a compositional system that says that meaningful wholes are built from meaningful parts guarantees that one perform any operations on those parts. Second, nothing about a compositional system guarantees that if ‘Peter loves Mary’ is well formed and meaningful that ‘Mary loves Peter’ is. We are merely told how the meanings of wholes depend on the meanings of parts. A perfectly good, compositional grammar can be constructed that allows ‘Peter’ to appear only in the verb phrase of sentences of the form ‘A loves B’. A compositional grammar as such simply isn’t the right kind of thing to account the richness of generality. So it would be a mistake to think that the LOT guarantees generality of thought.

But recall that generality is conceived of here as a *constraint*, so some simple backwards engineering is one option for those who take generality seriously.¹⁹ A proponent of the LOT could argue that what the Generality Constraint shows us is that the grammar of mentalese is indeed itself general – it is not the sort of grammar that allows for ‘Peter loves Mary’ but not ‘Mary loves Peter’. In effect, generality is assumed and built into the grammar rather than falling out of the grammar. Furthermore, a LOT theorist could hold that in light of the Generality Constraint there must be computational operations on mental sentences that allow one to access mentalese words. I don’t see any reason that proponents of the LOT should be resistant to these ideas if they also take the Generality Constraint seriously.

¹⁹ See fn. 13 above.

But even with those clarifications in mind, the LOT is still inadequate. Heck (2007, 2012) provides reasons, not only for thinking that the LOT comes up short, but also for thinking that it really is *content* in particular that should carry the burden of meeting the Generality Constraint.

Heck resists all manner of psychological explanation in terms of the LOT because they don't think such explanations would be sufficiently abstract – that is, they don't think such explanations can yield law-like, psychological generalisations. One law-like psychological feature is the fact that (all else equal) if a thinker can think the thoughts that a is F and that b is G, then she can think the thoughts b is F and that a is G. Heck's concern is brought out across two papers – Heck (2007), which concerns the Generality Constraint directly and Heck (2012), which concerns Frege's puzzle.²⁰ The point is put most clearly in the 2012 paper in a discussion of how we might explain the behaviour of a subject (Fred) who finds himself in a Frege-puzzle. Heck entertains a Russellian theory of content paired with a LOT and explains why explanations that make reference to LOT sentences won't allow us to offer a satisfactory psychological theory:

The explanation makes explicit reference to particular Mentalese sentences [...] and therefore lacks anything like the generality psychological laws formulated in terms of content were supposed to have. One can see this by considering what the corresponding psychological law would be: It too would make explicit reference to particular Mentalese sentences, and, as a matter of empirical fact, it would therefore probably have no other instances than Fred and so would lack the sort of generality we ordinarily suppose psychological laws to have. If everyone had the same language of thought – if the same sentence expressed the same content in my language of thought that it did in

²⁰ Thanks are due to a helpful conversation with Professor Heck, especially for pointing me to their 2012 paper and tying the papers together. Many of the underlying ideas in the 2007 paper are brought out in much more detail in the 2012 paper.

everyone else's – then that would be different, but we have no reason to suppose that is true. Indeed, there is no obvious reason to suppose that the notion of 'same Mentalese sentence' so much as makes sense interpersonally. It is for largely this reason that, even if one does accept the language of thought hypothesis, psychological laws must still be stated at the level of content. The computational story in which Mentalese appears is supposed to be a story about how psychological states and processes are implemented. It follows that explicit reference to sentences of Mentalese should no more appear in psychological laws than does explicit reference to neurons. (p. 152-3).

A theory of thought aimed at capturing the fact that thinkers who can think that a is F and can think that b is G can thereby think that a is G and that b is F is simply a special case. We take it as a psychological fact that, in normal circumstances, subjects who think certain thoughts are thereby capable of thinking other thoughts. Appeals to LOT sentences would rob us of a theory that achieves sufficient abstraction.²¹ But we get more from Heck, for they also tell us why appealing to content in particular is desirable. Doing so allows us to offer psychological regularities such as thought's generality in terms that *are* sufficiently abstract. Positing a LOT is of course compatible with this line, for LOT sentences have semantic values, but in order to offer psychological explanation that is law-like across individuals who may have, structurally speaking, extremely different brains, Heck holds that we must appeal not to syntactic facts or their kin, but to semantic facts. That is, *psychological explanation is intentional explanation*. Fodor seems to agree:

What's common to minds that have the same psychology is the intentional laws that subsume them, not either the computational or the neurological mechanisms that implement the intentional laws. Intentional states and

²¹ See also Evans (1982, p. 101). Evan's remarks are less specific, but he too shies away from appealing to a LOT and apparently for similar reasons to those given by Heck.

processes are multiply realized by computational states; computational states and processes are multiply realized by neurological states (or whatever), and, for all I know, neurological states are multiply realized by biochemical states; and so on down to (but not including) basic physics. For all I know, it can be multiply realized implementations all the way down so long as each of the multiple realizations implements the same intentional laws. (LOT2, p. 91)

In light of these considerations, it looks tempting to return to a view that makes sense of generality in terms of contents or, at any rate, in semantic terms. What we want is something law-like and abstract enough to capture the possible variation in implementation. But this looks to lead to the profligation of kinds of contents with the costs outlined earlier. But I think we can have it both ways.

Thought and Thinking-of

A theory of thought can meet the Generality Constraint in a sufficiently abstract way without positing distinct kinds of content. Generality requires, as Evan's argues, a series of abilities. It's natural to think that one who is able think the thoughts that a is F and that b is G is also able to think the thoughts that b is F and that a is G *because* that thinker can think of a, think of b, think of F, think of G, and combine the relevant elements in a general way. I wish to take such abilities seriously in their own right – thinking-of is itself a folk-psychological mental state and it features in a story of propositional thought. Cognitive states with propositional contents decompose into simpler states.

I do not hold that thinking-of is a mere abstraction over propositional attitudes. One who does may claim that one thinks of a just in case one forms some or other propositional

thought concerning a.²² Similarly, the proposal might continue, for thinking of F: One thinks of F just in case one forms a thought which predicates F of something or other. I reject this suggestion. Rather, on my proposal, thinking-of is a sub-propositional intentional state that is of or about things not in virtue of relating a subject to a proposition(s) concerning those things.^{23 24}

Since sub-propositional, intentional states have only recently begun to garner sufficient attention in the literature, allow me to say a bit more. Sub-propositional states such as thinking of the number seven, fearing snakes, liking John, and so on, much like the more familiar propositional attitudes, are *intentional*. When one thinks of, say, Mary, her thought is *about* or *directed upon* Mary; when one thinks of the property of being tall, her thought is about or directed upon that property. One crucial difference between propositional states such as belief and sub-propositional states such as thinking-of lies in predication. In thinking of o, nothing is predicated of o, but in a propositional thought about o, some or other property is predicated of o. Suppose I ask you to think of a number. You may think of, say, the number seven without predicating anything of it, you simply bring seven to mind. But if you believe that seven is prime, you predicate primeness of seven and hence take seven to be prime.

An important mark of a state's being intentional is the ability for it to be about things that don't exist. Sub-propositional states meet this mark; I can think of Pegasus or

²² Prior (1971) explicitly holds such a position and the view seems to be at least tacitly endorsed more broadly. One way to proceed is to take ascriptions of thinking-of o as shorthand for ascriptions of thinking-that o is F (for some appropriate F) or perhaps to take thinking-of o as a determinable of which all determinants are some or other propositional attitudes concerning o. See also Sainsbury (2010). For dissent and further discussion, see Crane (2001).

²³ See Ben-Yami (1997), Grzankowski (2012, 2016, 2018 a,b), and Montague (2007) for further development of attitudes towards objects. For more general discussion of non-propositional intentionality, see the essays in Grzankowski and Montague (2018).

²⁴ Thank you to Dr. sc. Pavel Gregorić for pointing me to *De Anima* III, 6 where Aristotle discusses thought and unity. By my reading, Aristotle appears to make space for separation in thought – separating various objects and properties and recombining them, especially important in false belief. He seems to hold that unities (what we might think of as whole thoughts) are prior to what's potentially decomposed for later combination. I can't pretend to be an Aristotle scholar (very far from it), but it interesting to see this early discussion of what looks to be closely related issues.

Vulcan for example. How to deal with states that are about things that don't exist is a familiar problem and as with propositional attitudes about things that don't exist, there are a range of ways one might proceed.²⁵ One might offer a view involving relations to abstract objects such as concepts or intensional objects. One might suggest that the states are relations to both existent and non-existent objects or perhaps to both actual and possible objects. Alternatively, one might wish to give up the relationality of intentional states as Uriah Kreigel (2011) and Angela Mendelovici (2018) have recently suggested. Given present interests and given that the same issues arise for propositional attitudes, I'll offer only the essential outlines of the kind of view I favour.

I suggest extending a broadly representationalist picture of the mind according to which thinking-of is analysed in terms of bearing a relation to an abstract, shareable, representational entity. In the case of propositional attitudes, representationalists appeal to relations to propositions. In the case of thinking-of, I suggest that we follow this model with only slight modification. In thinking of *o*, one stands in a relation to a concept that is about *o*. When *o* exists, the representation *refers* to *o*. When *o* does not exist, the representation cannot refer and is 'empty'. Since shareable concepts can be individuated other than by their referents,²⁶ empty concepts pose no special problem.

Returning to propositional thought, let us take *entertaining* a proposition as a base case. Entertaining a proposition, say that *a* is *F*, is a complex act of bringing objects and properties before the mind and taking the property to be instantiated by the object. The most basic propositional thoughts require three acts. When one entertains the proposition that *a* is *F*, one thinks of *a*, thinks of *F*, and one takes *F* to apply to *a*.²⁷ More complex

²⁵ See Grzankowski (2013) for further discussion.

²⁶ For example, one might appeal to conceptual role as in Block (1986) and Harman (1982) or to origin as in Sainsbury and Tye (2012).

²⁷ This observation makes important points of contact with Hanks (2011, 2015) and Soames (2010, 2015). According to Hanks and Soames, propositions are complex action types, so the proposition that *a* is *F* is the action type of predicating *F* of *a*. I reject this as a view of content (in part for reasons above concerning multiplying kinds of contents, but for further reasons still – see Grzankowski and Buchanan 2019) but there is much else in their picture of mental activity and predication that is very attractive.

thoughts are, accordingly, more complex: to entertain that aRb one must think of a , R , and b and take relation R to hold between a and b . Quantificational thoughts can also be accommodated by adopting the familiar idea from the theory of Generalized Quantifiers that express properties of properties.²⁸ To think the thought that every student sleeps, one must think of the property of sleeping, think of the property of being instantiated by every student, and must take the property of sleeping to instantiate the property of being instantiated by every student.²⁹ To simplify discussion, focus will be on simple subject-predicate thoughts of the form $[a \text{ is } F]$.

‘Taking’ (as in, ‘taking F to apply to a ’) is a further mental act; it is an operation on abstract representations. To take F to apply to a upon thinking of a and thinking of F is to bring one’s a -representation and one’s F -representation into an appropriate relation. One may wish to think of this as concatenation or as the mental analog of predication. In the case of the Standard Theory, we saw a theorist like Peacocke building into the possession conditions for predicate concepts a general mechanism that allows for the combination of predicate concepts with object concepts. On the present proposal, the relation of taking plays this role.

All propositional attitudes require entertaining a proposition on the present view.³⁰ When one judges that a is F , one entertains that a is F and endorses the application of F to a . When one believes that a is F , one entertains that a is F and is disposed to judge that a is F . For simplicity, I’ll continue to focus on entertaining.

²⁸ See especially Barwise and Cooper (1981).

²⁹ In many cases this isn’t something one does knowingly and consciously, but much as someone might hold that linguistic abilities rely on complex manipulations of deep syntactic structures that we aren’t ourselves aware of relying on, I am suggesting that complex instances of thinking propositionally is richer under the hood than what’s given to us in introspection.

³⁰ Again, this thought makes contact with recent work by Soames (2010, 2015). As above, Soames is offering a theory of propositions that I don’t myself endorse, but he helpfully suggests entertaining as a base case for propositional attitudes and then offers attractive ways of thinking of judgement, belief, and desire in terms of entertaining.

Returning now to the Generality Constraint, it is taken as a data point that if one can entertain that a is F and can entertain that b is G, then one can entertain that a is G and that b is F. On the present view, this is captured in terms the structure of thought in the following way: The thought that a is F requires thinking of a, thinking of F, and taking F to apply to a. The ability to entertain guarantees the representational sub-states and the ability to combine representations. Similarly for entertaining that b is G. But since having the ability to be in the pair of propositional states includes in it the abilities to be in the non-propositional states of thinking of a, b, F, G and the ability to relate representations, one is guaranteed to have all that's needed to form further thoughts. In particular, one has all that's needed to entertain that a is G and that b is F.

Structurally this approach is very similar to the Standard Theory in that it offers an account of what it is to be in a propositional attitude state in terms of a certain sub-structure. As I noted earlier, I agree with much in the Standard Theory. But notice that *content* is no longer carrying explanatory weight as I've spelled things out. Whereas the standard theorists build into content the concepts that must be grasped in order to have a contentful thought, the present view meets the Generality Constraint in terms of complex *states* and *acts* – no need to bring content into the story.³¹ So, when we turn to states that fail to meet the Generality Constraint we no longer reach the conclusion that they have a different kind of content. Rather, failures of generality will be due to either a failure of decomposition into sub-states, a failure to decompose into thinking-of, or a failure to demand predication. One might very reasonably, for example, maintain that one doesn't find generality in perception because in perceiving that a is F, one needn't think of a or think of F – perception involves no thinking. But granting this, one can still maintain that the contents of both perceptual states and cognitive states are of a kind.

³¹ An act that may endure. Forbes (2010) has recently suggested taking propositional attitude *states* to be enduring events. That thought meshes well with the present proposal.

The present explanation of generality also succeeds where the LOT did not. Rather than appealing to features of individuals that may well not generalise to other individuals, the explanation now on offer appeals only to psychological state types and acts of combination that might be realised in a wide variety of ways. A LOT remains an option at the level of implementation, but the description of thought presently offered appropriately abstracts away. Heck rightly demands that we account for generality and other psychological phenomena in a sufficiently general, law-like way. Their suggestion was that an appeal to content was hence needed. I don't want to downplay the role of content in various kinds of psychological explanation (for example, in belief/desire explanations of action), but when it comes to meeting the Generality Constraint, I think we can offer a perfectly acceptable theory in terms of state types and acts that makes no reference to the brain states of specific subjects or the token states of this creature or that. It is because of the constituent structure of entertaining propositions that the Generality Constraint is met: they are complex states that comprise other states. It is the existence of sub-propositional states that are part of folk-psychology that makes room for this story.

Conclusion

My primary aim has been to diffuse an argument for content nonconceptualism by providing an account of the structure of thought in terms of states rather than their contents. I argued that this was an important argument to diffuse since positing distinct kinds of contents comes at the cost of complicating the kinds of roles contents were meant to play. Content nonconceptualists aim to use the Generality Constraint to motivate the view that the contents of thoughts have concepts as constituents, but I've offered an alternative in terms of sub-propositional states that meets the Generality Constraint and avoids the kinds of objections leveled against the LOT all while remaining silent on the nature of content.

References

- Alston, W. P. (1998). "Perception and Conception." In *Pragmatism, Reasons and Norms: A Realistic Assessment*: 59-87. New York, Fordham University Press.
- Barwise, J. and Cooper, R. (1981). Generalized quantifiers and natural language. *Linguistics and Philosophy* 4 (2):159-219.
- Beck, J. (2012). The Generality Constraint and the Structure of Thought. *Mind* 121 (483):563-600.
- Ben-Yami, H. (1997). 'Against characterizing mental states as propositional attitudes.' *Philosophical Quarterly* 47 (186):84-89.
- Block, N. (1996). "Mental paint and mental latex." *Philosophical issues* 7: 19-49.
- Block, N. (1986). 'Advertisement for a Semantics for Psychology'. *Midwest Studies in Philosophy*, 10: 615–678.
- Brewer, B. (2011). *Perception and Its Objects*. Oxford University Press.
- Byrne, A. (2001). "Intentionalism defended." *The Philosophical Review* 110(2): 199.
- Camp, E. (2007). 'Thinking with maps'. *Philosophical Perspectives* 21 (1):145–182.
- Camp, E. (2004). 'The generality constraint and categorial restrictions'. *Philosophical Quarterly* 54 (215):209–231.
- Crane, T. (2001). *Elements of Mind*. Oxford, Oxford University Press.
- Davies, M. (1992). Aunty's own argument for the language of thought. In Jes Ezquerro (ed.), *Cognition, Semantics and Philosophy*. Kluwer. 235--271.
- Dickie, I. (2010). 'The generality of particular thought'. *Philosophical Quarterly* 60 (240):508-531.
- Evans, G. (1982). *The Varieties of Reference*. Oxford, Oxford University Press.
- Field, H. (1978). 'Mental representation'. *Erkenntnis* 13 (July): 9-18.
- Fodor, J. (2008). *Lot 2: The Language of Thought Revisited*. Oxford University Press.

- Fodor, J. (1998). *Concepts: Where Cognitive Science Went Wrong*. Oxford, Oxford University Press.
- Fodor, J. (1975). *The Language of Thought*. Cambridge: Harvard University Press.
- Forbes, G. (2010). 'Intensional verbs in event semantics'. *Synthese* 176: 227-242.
- Grzankowski, Alex (2012). Not All Attitudes are Propositional. *European Journal of Philosophy* (3):374-391.
- Grzankowski, Alex (2013). Non-Propositional Attitudes. *Philosophy Compass* 8 (12):1123-1137.
- Grzankowski, Alex (2015). Pictures Have Propositional Content. *Review of Philosophy and Psychology* 6 (1):151-163.
- Grzankowski, Alex (2016). Attitudes Towards Objects. *Noûs* 50 (2):314-328.
- Grzankowski, Alex (2018a). A relational theory of non-propositional attitudes. In Alex Grzankowski & Michelle Montague (eds.), *Non-Propositional Intentionality*. Oxford University Press.
- Grzankowski, Alex (2018b). Non-propositional Contents and How to Find Them. *Journal of Consciousness Studies* 25 (3-4):233-241.
- Grzankowski, Alex & Buchanan, Ray (2019). Propositions on the cheap. *Philosophical Studies* 176 (12):3159-3178.
- Grzankowski, Alex & Montague, Michelle (eds.) (2018). *Non-Propositional Intentionality*. Oxford University Press.
- Hanks, P. (2011). "Structured Propositions as Types," *Mind* 120, 11-52.
- Hanks, P. (2015). *Propositional content*. Oxford: Oxford University Press.
- Harman, G. (1990). "The intrinsic quality of experience." *Philosophical perspectives* 4: 31-52.
- Harman, G. (1982). 'Conceptual role semantics'. *Notre Dame Journal of Formal Logic* 28 (April):242-56.

- Heck, R. K. (2007). 'Are there different kinds of content?'. In B. McLaughlin and J. Cohen (eds.), *Contemporary Debates in Philosophy of Mind*. Malden, Blackwell.
- Heck, R. K., (2012). Solving Frege's puzzle. *Journal of Philosophy* 109 (1-2):132-174.
- King, J. (2007). *The Nature and Structure of Content*. Oxford University Press.
- Kriegel, U. (2011). *The Sources of Intentionality*. Oxford: Oxford University Press.
- Loar, B. (1988), 'Social Content and Psychological Content', in R. Grimm and D. Merrill (eds.) *Contents of Thought: Proceedings of the 1985 Oberlin Colloquium in Philosophy*. Tucson: University of Arizona Press.
- Lycan, W. G. (1996). *Consciousness and experience*, The MIT Press.
- McDowell, J. (1994). *Mind and World*. Cambridge, Harvard University Press.
- McGinn, C. (1989). *Mental Content*, Blackwell.
- Mendelovici, Angela (2018). *The Phenomenal Basis of Intentionality*. New York, USA: Oxford University Press.
- Montague, M. (2007). 'Against Propositionalism'. *Noûs* 41 (3):503–518.
- Pautz, A. (2010). 'Why explain visual experience in terms of content?' In Bence Nanay (ed.), *Perceiving the World*. Oxford University Press. pp. 254--309.
- Pautz, A. (2009). "What are the Contents of Experiences?" *The Philosophical Quarterly* 59(236): 483-507.
- Peacocke, C. (1992). *A Study of Concepts*. Cambridge, The MIT Press.
- Peacocke, C. (1983). *Sense and Content*, Clarendon Press.
- Prior, A.N. (1971). *Objects of Thought*. Oxford, Clarendon Press.
- Sainsbury, R.M. and Tye, M. (2012). *Seven Puzzles of Thought: And How to Solve Them: An Originalist Theory of Concepts*. Oxford, Oxford University Press.
- Schellenberg, S. (2018). *The Unity of Perception: Content, Consciousness, Evidence*. Oxford: Oxford University Press.
- Schiffer, S. (2003). *The Things We Mean*. New York: Oxford University Press.

- Searle, J. R. (1983). *Intentionality, an essay in the philosophy of mind*, Cambridge University Press.
- Shoemaker, S. (1990). "Qualities and Qualia: What's in the Mind?" *Philosophy and Phenomenological Research* **50**: 109-131.
- Siegel, S. (2011). "The Contents of Perception", *The Stanford Encyclopedia of Philosophy*, Edward N. Zalta (ed.),
<<http://plato.stanford.edu/archives/win2011/entries/perception-contents/>>.
- Siegel, S. (2010). Do visual experiences have contents? *Perceiving the World*. B. Nany, Oxford University Press.
- Soames, S. (2010). *What is Meaning?* Princeton, Princeton University Press.
- Soames, S. (2015). Rethinking language, mind, and meaning. Princeton: Princeton University Press.
- Stalnaker, R. (1990). Narrow content. In C. Anthony Anderson & Joseph Owens (eds.), *Propositional Attitudes: The Role of Content in Logic, Language, and Mind*. Stanford: Csl.
- Thau, M. (2002). *Consciousness and cognition*, Oxford University Press, USA.
- Travis, C. (2004). "The silence of the senses." *Mind* **113**(449): 57.
- Tye, M. (2000). *Consciousness, Color, and Content*. The MIT Press.