Finite thinkers

Olivia Sultanescu

Concordia University

Invited chapter for *Introducing Philosophy of Mind, Today* (Routledge)

Editors: Devin Curry and Louise Daoust

> We must allow ourselves to think, we must dare to think, even though we fail. It is in the nature of things that we always fail, because we suddenly find it impossible to order our thoughts, because the process of thinking requires us to consider every thought there is, every possible thought. (Thomas Bernhard, *Extinction: A Novel*)

Because we are thinking creatures, understanding ourselves requires understanding what it is to have thoughts. Thoughts are unlike ordinary material items, such as tomatoes. We can hold tomatoes in our hands, we can smell and taste them, and we can measure their proportions. But we can do none of these things with thoughts, for thoughts are not observable, at least not in the way in which tomatoes are. In fact, the very idea that thoughts are individual things should strike us as puzzling. How many thoughts about my cat do I have, or have I had, over the course of my life? And is the thought that 1+3=4 the same as or different from the thought that 3+1=4? We seem to have a straightforward conception of how to count tomatoes, but we are at a loss when it comes to counting thoughts.

One thing is clear: the only way to single out thoughts is by attending to what they are about, or, as philosophers put it, by attending to their intentional content. Thoughts are directed toward objects; they are *about* bits of reality. For instance, my thought that Iris Murdoch is a philosopher is about Iris Murdoch. But this is not all there is to my thought's aboutness or

intentionality. My mind is fixed on Iris Murdoch in a distinctive way, for I take her to have a specific feature or property, namely, being a philosopher. My thought is true if she has that property and false otherwise.

Aboutness sets thought apart not only from tomatoes, which cannot intelligibly be said to be *about* anything, but also from other sorts of representation. Perhaps any representation of a thing somehow depicts it as exhibiting some property. But not every representation is true or false. A photograph is a representation, but a photograph cannot be said to be true or false in any strict sense. Photographs may be accurate or inaccurate, but accuracy, unlike truth, comes in degrees. Perhaps our perceptions are representations of the world. And yet, whether perceptions can be said to be true or false is not something that is immediately settled by this characterization, because the kind of content that perceptions have might not be the sort that renders them true or false. Indeed, perceptions may be more like photographs than thoughts in this respect. Thus, we shall provisionally say that a distinctive characteristic of thoughts is that they are true or false. We shall return to this issue.

What is it for thoughts to have content of the sort that renders them true or false? Given that the individuation of thoughts is in terms of their content, this is tantamount to asking what it is for thoughts to be the thoughts that they are. This is a constitutive question, for it seeks to unearth the nature or constitution of a phenomenon. The questions asked by Socrates in Plato's dialogues often take this form. As those dialogues also show, it is not easy even to understand such a question, let alone answer it. In the first part of this essay, I shall try to clarify the constitutive question about thought by further examining the characteristic of thoughts mentioned earlier, namely, their being true or false. In the second part of the essay, I shall consider an approach that is often taken to be best suited for addressing the question, and I shall

try to show that it cannot address it. This will allow us to articulate a puzzle, which seems to me to be central for our understanding of ourselves as minded beings bound to live finite lives. As we shall see, our finitude is not something that can be set aside for the purposes of the philosophical inquiry into the mind; grappling with it is an essential component of this inquiry.

The structure of thought

Conceiving of thought as having content would seem to invite us to think of the mind as a container, and of the contents of thoughts as elements that fill that container in the way in which tomatoes fill a bucket. But, as we suggested earlier, the contents of thoughts are special, for they are not mere collections of items. Thoughts exhibit structure: one element of the structure picks out an object; the other expresses a property ascribed to or predicated of that object. These two elements must be combined in the right way for them successfully to express anything. Thus, I cannot think that *Murdoch a philosopher Iris is*, because there is, in this case, nothing to be thought, nothing that can count as true or false. A thought is true if and only if the element picked out has the property ascribed to it. My thought that Iris Murdoch is a philosopher is true if and only if she is a philosopher. This structure, which philosophers may call 'predicative', is not the only one that thoughts can have, but it is both basic and pervasive, and our focus will be on it. From now on, whenever we speak of content, we will be referring to the sort of content that can be captured by a declarative sentence.

We should note, however, that some thoughts that may appear to have a predicative structure do not seem to pick out any real entity. The thought that the current president of Canada is a philosopher belongs to this category, for Canada is led by a prime minister, not a president. Moreover, we can have thoughts about things that we recognize to lack reality, such

as fictional characters. This raises the question of how it is possible to think about things that do not exist. And are thoughts about things that do not exist true or false, or are they neither? In what follows, we shall ignore these complications and focus instead on the arguably more fundamental question of how it is possible to think about things that exist. Thoughts about things that exist (or about things that have existed, such as the thought that Iris Murdoch is a philosopher) are indisputably true or false.

I cannot think of Iris Murdoch that she is *a philosopher* unless I can ascribe this property to other things, for instance, to other people, to aliens, to orange cats, etc. A bit of the world can be said to be a particular way only if other bits of the world can be said to be that way. And, arguably, I cannot think *of Iris Murdoch* that she is a philosopher unless I can ascribe other properties to her, such as being a writer, being a human being, being a bit strange, or being a spatiotemporal entity. A bit of the world can be said to be a particular way only if it can also be said to be other ways. This reveals another respect in which thoughts are unlike tomatoes: one can hold only one tomato in one's hand, but one cannot hold only one thought in one's mind.

Concepts are components of contents. Roughly, they are ways in which various bits of the world might be said to be. It is in the nature of a concept that it captures some bits of the world correctly and other bits incorrectly; it is thus in the nature of the concept *philosopher* that it applies correctly to some items and incorrectly to others. The concept applies correctly to Iris Murdoch, and incorrectly to the presidential candidates for the 2024 election in the United States. The distinction between correctness and incorrectness is essential to concepts, just as the distinction between truth and falsity is essential to thoughts.

Even if there was only one philosopher in the actual world, and thus only one object that could correctly be subsumed under the concept, the concept could still be incorrectly applied to

indefinitely many actual objects. And even if there were only one actual object, it is still conceivable that the concept could apply to indefinitely many objects, in part because it is conceivable that the world could be a different way than it is. This is what distinguishes a concept from a mere tag or label, a predicate from a mere name. Concepts exhibit a distinctive kind of generality, captured by the possibility of indefinitely many objects being subsumed, correctly or incorrectly, under them. To put it differently, indefinitely many thoughts may be formed with each of them. For instance, the concept *philosopher* can feature in indefinitely many thoughts, not just about what there is, but also about what there is not, about what could and could not be. This is an essential feature of concepts, and thus of thoughts, to which we shall return. Let us state it again below:

> The **generality of thought** consists in the fact that the concepts deployed in thought are such that they can intelligibly be applied, correctly or incorrectly, to indefinitely many objects.

What counts as a thought?

What we have said so far might be taken to suggest that the umbrella of thought encompasses only beliefs or convictions, cases in which the thinker takes something to be true. But the predicative structure we just described is distinctive of a much wider range of mental attitudes. This allows us to explain crucial commonalities between them. When I *hope* that Iris Murdoch is a philosopher (which I might do as I am about to begin reading an essay of hers for the first time), the content of my hope is the same as the content of the conviction that she is a philosopher. The same holds with respect to my (heaven forbid!) *doubting* that she is a

philosopher. More generally, it seems that there is a wide range of attitudes that I can take toward Murdoch's being a philosopher: I can wish for it, fear it, expect it, regret it, and so on.

Thus, if we take predicative structure to be characteristic of thought, we must allow that we are operating with a more capacious notion of thought than we initially suggested. Any mental attitude whose content exhibits that structure counts as a thought. It follows, then, that, strictly speaking, and contrary to the provisional characterization we offered earlier, only some thoughts are true or false, namely, those thoughts that are beliefs, convictions, and so on. But, in so far as all thoughts involve concepts—representations of ways in which the world might be—all thoughts exhibit the generality we described in the previous section.

The constitutive question

Now that we have circumscribed our topic, let us return to our question: What is it for thoughts to have content? What this question calls for is not an investigation of our ways of thinking, but rather an investigation of what makes it the case that we think in the ways that we do. Since for thoughts to have contents is, in part, for them to involve concepts, we are interested in what makes it the case that we have the concepts that we do.

In order further to spell out this constitutive question, let us consider a different, simpler, scenario. Take my concept of *red*. Suppose that you show me a tomato, one which I have not seen before. What makes it the case that the concept is correctly applicable to it? You might think that it depends on whether the tomato is red; if the concept is correctly applicable to red things, then it is correctly applicable to the tomato (if it is indeed red). But the question we are asking is more fundamental: What makes it the case that the concept is correctly applicable to

red things in the first place? This is tantamount to asking what makes it the case that I have the concept of *red* (as opposed to some other concept or no concept at all).

It is tempting to think that, in response to this question, we could point to some features of my human psychology, and hope to arrive through empirical inquiry into our cognitive mechanisms at a sufficiently detailed specification of those features. Presumably, my having the concept of red has to do with my belonging to a species that is capable of colour vision, endowed with certain skills, located in a certain kind of environment, responding to certain kinds of stimuli, and so on. But there is a deep difficulty that threatens such an approach. In a nutshell, it is hard to see how we can be thinkers at all, given the apparent tension between the finiteness of our nature and the generality distinctive of thought. Let us say more about the approach, which we shall call *dispositionalist*, in order properly to articulate the tension.

The dispositionalist approach

Our human nature, which is supposed to enable us to specify what makes it the case that we have the concepts we do, can be described in terms of a complex configuration of dispositions, or tendencies to behave and react in particular ways. And if we think of our nature in this way, it might seem that we can explain what it is for an individual's thoughts to have contents, and thus answer the constitutive question, in terms of the individual's having certain kinds of dispositions.

In fact, this sort of explanation seems especially promising, given the feature we identified as essential to thought, namely, its generality. Dispositions are the sorts of things that can be manifested in a wide, potentially infinite, range of situations. For example, if we consider the disposition of a glass to break, it might seem that there is an infinite number of circumstances that can trigger it. There are, for instance, infinitely many scenarios involving a table the glass

might fall off. So, there is a sense in which the disposition of the glass has an infinitary character. If a glass has infinitary dispositions, why could we not view the human being as having some as well? And if we can indeed view her in this way, there might be no obstacle to accounting for the generality of thought in terms of dispositions.

There are competing conceptions of dispositions endorsed by philosophers who focus on them in grappling with the constitutive question, but we need not explore them here. This is because we need not adjudicate the question of which one of them is right. What we are concerned with is the feasibility of the basic idea behind this approach. And what the basic idea amounts to in our case is that to think that the tomato is red is to be disposed to behave in particular ways in response to the tomato—for instance, to be disposed to utter the word 'red' (or 'rouge', or 'rojo', or 'roşie', etc.) when asked about its color, to be disposed to place the tomato in the same category as stop signs, strawberries, and scarlet dresses, and so on. Crucially, the relevant dispositions must be able to be characterized in a way that does not rely on the idea of thought or content, on pain of rendering the view vulnerable to the charge of circularity, according to which one is taking for granted that which one is trying to explain.

The problem of finitude

Suppose that we grant that my possessing the concept of *red* consists in my being disposed to respond in certain ways to red things. And now consider this fact: all my past applications of the concept *red*, indeed, all my thoughts involving the concept, are, on the face of it, perfectly consistent with the following interpretation: the concept applies correctly both to red things encountered prior to July 14, 2024 and to blue things encountered starting July 14, 2024. All my past uses of the concept are consistent with this interpretation because, let us suppose, all my

encounters with the world took place before July 14, 2024. This interpretation preserves the truth (or falsity, as the case may be) of all my thoughts involving redness. Of course, it might strike us as implausible, even ludicrous; surely, nobody would ever be tempted to ascribe this bizarre concept to me. But the question is whether the philosopher who thinks that having a thought is merely a matter of manifesting certain dispositions has the philosophical resources to show that the interpretation is false. What makes it the case that the interpretation according to which the concept applies correctly to all red things, even red things after July 14, 2024, is the right one?

One might think that an answer is available on a dispositionalist picture, given that there is a fact of the matter about what my dispositions are around (and after) July 14, 2024. But nothing prevents us from devising an interpretation that cannot be so easily ruled out. Consider the following suggestion: my concept *red* applies correctly to red things encountered prior to July 14, 3024, and to blue things encountered starting July 14, 3024. Is there a fact of the matter that can rule out this interpretation? I will no longer be around one thousand years from now, and so I cannot be said to have any dispositions covering this future segment of time.

Crucially, despite the absence of the relevant sort of dispositions, the concept that I am using now, at this moment, is such that it does not exhibit the strange characteristic of picking out blue things after July 14, 3024. It is a feature of the content of my present thoughts about red things that the concept deployed in those thoughts picks out red things regardless of time and place. Things that one might encounter one thousand years from now properly belong under my present concept if they are red. The concept somehow reaches out into the future in this way, a way that is not limited by the circumstances of my actual life. This is indeed what the generality

of thought amounts to, namely, the fact that the elements of a thought somehow contain infinitely many possible thoughts. How can we make sense of my thoughts possessing this feature?

Here is another suggestion: instead of attending to my actual dispositions and their usual circumstances, we should consider them against the background of idealized circumstances, which have yet to be specified. An adequate idealization, the thought goes, would enable us to rule out bizarre interpretations, such as the one we just proposed, for it would allow us to say that I do not lack the relevant dispositions. The impression that I lack them is due entirely to contingent limitations that the philosopher can abstract away from. To put the point differently, what is required is an uncoupling of my dispositions from the circumstances of my *actual* life, followed by a consideration of them in *ideal* circumstances. What might these circumstances be?

Obviously, the line of thought continues, it would be ideal if I lived an indefinitely long life, in no small part because this would allow us to say that, as a matter of fact, I have the dispositions that I was thought to lack. For instance, I do have a disposition concerning red things encountered starting July 14, 3024. On this picture, to have the concept of *red* is to be disposed, in circumstances in which one lives an indefinitely long life, to apply the concept to red things. More generally, having a concept is indeed a matter of having a disposition, with the caveat that this is taken to incorporate an appeal to idealized circumstances. This can provide us with the skeleton of an attractive answer to the constitutive question, or so it might be surmised.

The trouble is that idealizing circumstances in this way is a questionable move. I am a human being (alas!), and thus a finite biological creature, who is alive for a limited period of time. Due to advances in medical technology, human beings may live increasingly long lives. But their lives will never cease to be finite. Thus, trying to picture myself as someone who can live an indefinitely long life is not a case of supposing myself, for the sake of theorizing, to have

a feature that I in fact lack; paradoxically, it guarantees that I am no longer on the scene, in so far as it shifts the focus onto a fundamentally different kind of being. To put the point differently, what we need to explain, what the constitutive question we are concerned with ultimately asks, is what makes it the case that a finite being like me is capable of the sort of generality distinctive of thought. Shedding light on how a being who is not finite might be capable of such a feat amounts to changing the subject.

The deep difficulty in the attempt to account for the generality of thought in dispositionalist terms that we just uncovered is a theme in the work of both Ludwig Wittgenstein and Saul Kripke (Wittgenstein 1953, Kripke 1982). I believe that this difficulty is insufficiently appreciated by contemporary philosophers. It certainly poses a challenge to dominant views. How shall we proceed, now that we acquired an appreciation of it? We seem to be faced with two options: we can either deny the generality of thought, or we can opt for a characterization of human nature that presupposes it. The latter path requires that we reconsider our explanatory ambitions. Let us examine each path in turn.

Rejecting generality

You might be wondering what is so unappealing about the claim that our concepts do not exhibit the sort of generality that we ascribed to them. Perhaps concepts are such that they can intelligibly be applied only to things that are within human reach. For instance, perhaps my concept of *red* is such that it applies correctly only to things that are red either at this moment or in the not-too-distant future (or past). If there is a tomato on Mars a millennium from now (and supposing that no human beings will be around on Mars a millennium from now), why should we think that it is a feature of my current concept that it applies correctly to it?

To begin with, note that this question does not strike us as gibberish; it is a perfectly intelligible question. This shows that we are able to think and talk, by relying on our current concepts, about goings-on that take place in a world in which there are no humans. On the face of it, giving up on the generality of thought requires that we deny this. If our thoughts and concepts cannot reach beyond the domain delimited by our dispositions, then presumably we can think and talk only of that domain. The appearance that we can think about the future of Mars is, on this picture, just this: an appearance that is ultimately illusory.

The troubling nature of the move of rejecting generality as a feature of thought can also be brought out if we attend to the mathematical domain. Consider the following fact: there are numbers that are too large for human beings to grasp. Any human being would die before fully thinking them through. It seems indisputable that, just by drawing on our ordinary conception of numbers, we can make sense of there being such numbers. Now, suppose that I am competent with addition, such that I can have thoughts about sums. We can think of the concept of sum as a relational concept, which relates three numbers. It applies correctly to the sequence <2, 3, 5> (since two plus three is five) and incorrectly to the sequence <5, 5, 11> (since five plus five is ten, not eleven). Like any other concept, it also applies correctly or incorrectly to indefinitely many other sequences. Now, consider sequences composed of numbers that cannot be grasped by human beings. By definition, not only do I *not* have a disposition concerning such numbers, but *no human being*, past, present, or future, has the relevant disposition. Nevertheless, it is a basic feature of our ordinary understanding of arithmetic that the concept applies correctly to some sequences of such numbers and incorrectly to other such sequences, and that whether adding two numbers yields a determinate result is independent of our ability to grasp that result.

If we give up on this feature, we are effectively subordinating arithmetic to the feebleness of our contingent human nature.

This suggests that the rejection of generality comes with unacceptable costs. Are the costs associated with accepting generality less severe?

Reforming our ambitions

We suggested earlier that one option would be not to renounce generality but to conceive of our nature as already encompassing it. For instance, on this picture, to have the concept *red* is to be able to think of things that they are red. My possession of the concept is in this way presupposed in the characterization of my abilities. As a result, these abilities can no longer be seen as more basic than my possession of the concept. Is there anything that prevents us from taking this path?

Recall that our question is what it is for thoughts to have content, and that, given that the individuation of thoughts is in terms of their content, this is tantamount to asking what it is for thoughts to be what they are. In saying that to have the concept *red* is to be able to think of things that they are red, we seem to be proposing that for thoughts to be what they are is simply for them to be the thoughts that they are. But this account does nothing more than restate the very idea that it is supposed to explain. It is evidently circular.

A widespread assumption among philosophers is that we should expect constitutive questions to be able to receive answers that are not circular. Taking our human psychology to involve, at its core, abilities to think of the world in certain ways must come with the admission that the aspiration to articulate non-circular answers is misguided in the case of the constitutive question about thought. But does this mean that we must altogether abandon the attempt to

address constitutive questions about thought? Not necessarily. Following P. F. Strawson, we might distinguish two models of philosophical inquiry, and thus two ways of approaching constitutive questions (Strawson 1992, 19-20). On one model, which Strawson calls *reductive*, circularity is always viewed as a flaw. On an alternative model, which Strawson calls *connective*, circularity has the potential to be illuminating. If we subscribe to the connective model, our ambition is not to answer constitutive questions about a phenomenon in a non-circular manner, but rather to articulate the essential connections between that phenomenon and other phenomena that cannot themselves be reductively accounted for. Plausible candidates in our case are the connections between the ability to think and other abilities of ours, such as the ability to act, use language, and know oneself, others, and the world. None of these abilities can be illuminated except in the light of the others, but this does not mean that they cannot be illuminated at all.

The connective model, albeit much less popular than the reductive one, is not unexploited. One philosopher who developed a comprehensive conception that exemplifies this model is Donald Davidson. On this conception, the content of a thought "derives from its place in the pattern" of beliefs, preferences, intentions, hopes, fears, and expectations (Davidson 1970, 221). This pattern can be instantiated only in creatures who are believers of truths and lovers of the good. Luckily, our finitude notwithstanding, we are such creatures. We cannot say what it is for our thoughts to have the contents they do in terms that do not presuppose that we are thinkers, and thus that do not presuppose that we are inherently capable of the generality distinctive of thought. But we can try to articulate illuminating connections between our being thinkers and other aspects of our nature.

Conclusion

We are finite creatures, and to be a finite creature is to have a finite life. But we are also thinking creatures, and what it is to be a thinking creature is for one's mind to exhibit a distinctive variety of generality. These two aspects of our nature are in tension. But neither of them seems to be dispensable: we really are made of flesh and bones, and we really are capable of thought. How can this be? Understanding ourselves requires grappling with this puzzle, which seems to me to reveal something significant about the sort of creatures we are.

One way to deal with the puzzle is to deny that human minds are capable of the distinctive variety of generality that we have uncovered. Another way, which I take to be far more attractive, is to reconceive of our explanatory ambitions. We cannot offer a reductive explanation of how it is that, despite our finitude, we are capable of thought. But this need not mean that we can offer no philosophical illumination of thought.[*]

References

Davidson, Donald. 1970. "Mental Events." In *Essays on Actions and Events*. Oxford: Clarendon Press, 2001.

Kripke, Saul. 1982. *Wittgenstein on Rules and Private Language: An Elementary Exposition*. Oxford: Blackwell.

Strawson, P. F. 1992. *Analysis and Metaphysics: An Introduction to Philosophy*. New York: Oxford University Press.

[*] Thanks to Claudine Verheggen, Samuel Steadman, Ulf Hlobil, Nabeel Hamid, and Alex Miller for comments and discussion. I am also very grateful to the editors of the volume for many helpful suggestions.

Wittgenstein, Ludwig. 1953. *Philosophical Investigations*. Edited by P. M. S. Hacker and Joachim Schulte. Translated by G. E. M. Anscombe, P. M. S. Hacker, and Joachim Schulte. Fourth edition. Wiley-Blackwell.

Further readings on the finitude problem

Boghossian, Paul. "Is Determinate Meaning a Naturalistic Phenomenon?" In *Meaning Without Representation: Essays on Truth, Expression, Normativity, and Naturalism*, edited by Steven Gross, Nicholas Tebben, and Michael Williams, 2015.

Miller, Alexander. "Kripke's Wittgenstein's Sceptical Argument", in *Wittgenstein and the Possibility of Meaning: "To Follow a Rule Blindly."* Oxford, New York: Oxford University Press, 2024.