

Conceptual Analysis and its Limits¹

Karen Bennett
Cornell University

My topic is, well, conceptual analysis and its limits. I will start by sketching what I mean by ‘conceptual analysis’, and saying a bit about how it is used in contemporary philosophy. Then I will point out two limitations of the method, and illustrate these limits with examples: some from the philosophical literature, and some from biology.

1. Conceptual Analysis

Conceptual analysis is the use of thought experiments and reflection on possible cases to reveal to ourselves the boundaries of our concepts: what it is for something to count as an *F*. The idea is to try to bring the contours of an ordinary concept into sharper focus. We learn from our reactions to cases; we learn something about what we think could count as a thing of a certain kind (c.f. Jackson 1998, 31-32).

Here’s a banal example. Suppose you are offering an account of what it is for something to be a sandwich. You say: a sandwich is a food made from meat and/or cheese and/or vegetables enclosed in bread. But then I problematize your analysis by asking, “is a hot dog a sandwich?” My own gut reaction is no, it’s not a sandwich. If you agree, we need to go back and refine our first pass account. Maybe a sandwich is actually *sliced* meat and/or cheese and/or vegetables enclosed in bread. But what about egg salad and tuna salad? Our theory needs further refinement! At this point, however, I propose abandoning the inquiry; obviously philosophers don’t actually spend a lot of time developing careful theories of The Sandwich.

Now, that example might make you think that conceptual analysis is nothing but messing around with words. That’s not

¹ This was written as a “public lecture... intended for a general audience” at the Center for Philosophic Exchange at SUNY Brockport. Please read it in that spirit. Thanks to the Brockport audience for discussion.

exactly wrong: we're trying to figure out what exactly the word 'sandwich' means, what our concept of a sandwich is. And as I have already suggested, nobody much cares about *that* question. But there are more interesting cases. I shall quickly sketch a couple of familiar examples, both to further illustrate how conceptual analysis works, and to extract some thoughts about what conceptual analysis is used *for*.

First, knowledge. For a long time people thought that knowledge was justified true belief. That is, they thought that if you believe p on the basis of good reasons, and p is true, then you count as knowing p . But then Edmund Gettier came along (1963) with a kind of counterexample now known as a "Gettier case". Imagine that you're taking a walk on a friend's sheep farm. You see a sheep-shaped rock by a tree, and form the belief that there's a sheep by the tree. This is justified: after all, it looks like a sheep, and you do know that there are sheep around. And let's further suppose that it's *true* that there is a sheep by the tree: there's one just behind the rock that you've mistaken for a sheep. Do you *know* there's a sheep by the tree? It has seemed to many that no, you do not. You do not know that, even though you do have a justified true belief to that effect. So much for the analysis of knowledge as justified true belief.

Second, free will. Lots of people care very much about whether we ever act freely. But whether or not we do might well depend on what free action *is*. What does it take for an action to be free, or to be freely chosen? Here's one possibility: it requires that you could have chosen differently. You chose to read this paper, but you could have picked a different article, or chosen not to read a philosophy paper today at all. So your action is free because you could have acted differently. Now, if *that's* what it means to act freely, then it's tempting to think that your acting freely requires the falsity of determinism. (I only say that it is tempting. There are ways to try to claim that determinism is compatible with the ability to do otherwise; e.g. Lewis 1981). The idea of determinism is familiar: it says that events at time t are fully settled by the laws of nature and earlier events. There is only one possible outcome. So if acting freely requires the availability of genuine alternatives, then if determinism is true, we never act freely (or so the thought goes). But of course many people—compatibilists—think that free action is compatible with determinism, as long as you are careful about what it takes for an action to count as free. For example, Walter Stace (1952) defended the admirably simple view that a person acts

freely just in case their action is caused by their own internal psychological states—what they believe and want. Actions aren't free when they are externally compelled. On his view, there's a big difference between lying on the floor because you want to, and doing so because someone pushed you down. On *this* account of what free action is, it doesn't matter whether or not determinism is true.

Obviously I just opened a very large can of worms. So let me be clear that I'm not trying to do more than gesture towards the vast literature on free will. My only point is that whether or not we ever act freely depends, among other things, on what it would *mean* to act freely. What if any actions count as free? The ones such that we could have done otherwise? Or the ones that are caused by the agent's desires and choices? Or something else? And these are questions about a concept: *what is our concept of free action?*

Now, something that's interesting about these two cases is that they actually illustrate two different uses of conceptual analysis. One use is to counterexample some particular analysis of what it is to be an *F*; this is what is at issue in the Gettier example. There, a particular account of what counts as knowledge—justified true belief—is challenged by the fact that we don't think you actually know there's a sheep by the tree. Our reaction to that kind of case reveals that that the justified true belief analysis doesn't capture our concept of knowledge.

The other use is to help figure out whether or not there are any *F*s. This is what is at issue in the free will case. To settle whether the truth of determinism would or would not mean that we never act freely, we have to sort out what a free action would be. Let me further illustrate this with a simple example. Consider the question, "are there witches?" To make progress on this question, we need to know what we are asking. What is a witch supposed to be? What would count as a witch? If witches are women who can really make magic potions, cast real spells, fly on broomsticks, then there aren't any. But if witches can be women with nontraditional religious beliefs who like to spend time in the woods, and maybe know a lot about herbs and herbal medicine, then there are some.

Both uses have at their core the thought that we need to define the subject matter of an inquiry before we start. We need to know what we have in mind, what an *F* would be, in order to meaningfully pursue questions about *F*s. Is this analysis of *F*hood a good one? Are there any *F*s? Here's Frank Jackson, noted defender /proponent of conceptual analysis:

Although metaphysics is about what the world is like, the *questions* we ask when we do metaphysics are framed in a language, and thus we need to attend to what the users of the language mean by the words they employ to ask their questions. When bounty hunters go searching, they are searching for a person and not a handbill. But they will not get very far if they fail to attend to the representational properties of the handbill on the wanted person. These properties give them their target, or, if you like, define the subject of their search. Likewise, metaphysicians will not get very far with questions like: Are there *Ks*? Are *Ks* nothing over and above *Js*? and, Is the *K* way the world is fully determined by the *J* way the world is? in the absence of some conception of what counts as a *K*, and what counts as a *J* (Jackson 1998, 30-31).

This idea, generalized, leads to what has become known as the Canberra Plan.

The Canberra Plan is named after the town where the Australian National University is located, and it's named after that town as a kind of surrogate way of naming it after two people: Frank Jackson and David Lewis. Jackson spent much of his career at the ANU; Lewis went to Australia every summer, and was both influenced by and influential in the Australian philosophical community. The Canberra Plan is basically a form of two-stage inquiry:

Stage 1: do some conceptual analysis on whatever issue is in question: free will, personhood, pain.

Stage 2: look around the world to see what best fits the description yields at stage 1.

If nothing fits the description, then the thing you were investigating doesn't exist. Of course, there are complexities here involving cases in which more than one thing fits the description, or something only partly fits the description (Lewis 1972). But they won't really play a role in what's to come.

Stage 1 of the Canberra Plan is the conceptual analysis stage. The idea, again, is that conceptual analysis is important ground-clearing work that sets the stage for further investigation. And that

has basically been the point of this section—of the knowledge and free will examples, and my discussion of the uses of conceptual analysis.

2. The Limits of Conceptual Analysis

So far, so good. But conceptual analysis has its limits, and in the rest of this paper, I want to call attention to and illustrate two of them. But before even telling you what they are, I want to flag that I am far from the only person to raise concerns about conceptual analysis. For example, Laura Schroeter (2004) argues that our grasp of the content of our own concepts is fallible. And Sarah-Jane Leslie and Mark Johnston (2012) argue that paying attention both to the way psychologists think about concepts and to the nature of generic generalizations changes how conceptual analysis looks. I'm not going to engage with this other work here; instead I will just raise two issues that are independent of their concerns. Let me be clear: I'm not in fact a foe of conceptual analysis, and I am not trying to rid philosophy of it. I think conceptual analysis is an excellent tool as long as one remains explicitly aware of its limits and pitfalls. My task here is to call attention to exactly that.

Here are the two limitations to which I want to call attention. First, conceptual analysis by itself doesn't tell you what concepts we ought to be deploying. Second, conceptual analysis by itself doesn't tell you whether anything in the world answers to the concept. This second limitation amounts to saying that the second stage of the Canberra Plan must not get left by the wayside. Let me spell these out in turn, by means of some cases that I hope are independently interesting.

3. Limitation 1: Conceptual analysis by itself doesn't tell you what concepts we ought to be deploying.

Conceptual analysis is by its nature *descriptive*. It is a tool for figuring out the contours of our existing concept of a person, or free action, or whatever. It can say nothing about whether our existing concepts are good ones. That is, it offers no guidance at all on when we ought to *revise* our concepts; it offers nothing *prescriptive*.² And

² Compare Strawson's distinction in the second sentence of *Individuals*: "descriptive metaphysics is content to describe the

this matters, because there indeed are circumstances under which we ought to revise—or at least ought to consider revising—our concepts.

One such circumstance is when a concept turns out to be empty. I'll briefly come back to this at the end of the paper. Another such circumstance, the one I want to pursue now, is when we have 'folk' notions of stuff that scientists also tell us about—for example, we have folk notions of what it takes to be a planet, to be solid, to have a spatial location.

One case that I find interesting is that of an *individual organism*. Now, obviously most of us do not use the phrase 'individual organism' in casual conversation, but we do indeed have a basic, untutored, 'pre-theoretic' or 'folk' notion of this. We know how to individuate and count cats and dogs; we know the difference between one bush and three bushes; we know how many daffodils we planted by the driveway. This ability to individuate and count shows that we have the concept, whether or not we could ever articulate or define it.

But of course 'individual organism' is also a scientific concept, used by biologists. *And their use doesn't always match the folk use.* There are cases in which there are scientific reasons to count as one what the folk would count as many; there are cases where there are scientific reasons to count as many what the folk would count as one. (Or, perhaps better, there are scientific reasons to not use the folk notion of an individual organism, but to use multiple replacement notions, depending upon the theoretical context.) But let me quickly sketch a couple of examples.

The first example is of something that the folk treat as a single organism, but which biologists treat as many: the Portugese man-of-war. It is not a true jellyfish, but a siphonophore, and it is considered to be a colonial animal. It consists of four polyps working in concert; it is more like a (small) beehive than a single animal. The details of why biologists consider it to be a colonial animal do not matter. What matters is that the decision to consider it not an 'it' but a 'they' belongs to biologists, not the folk.³

actual structure of our thought about the world, [sic] revisionary metaphysics is concerned to produce a better structure" (1959, 9).

³ For a little more information, see <http://www.siphonophores.org/index.php> and <http://animals.nationalgeographic.com/animals/invertebrates/portuguese-man-of-war/>

The second kind of case involves things the folk treat as many, but there is at least some scientific reason to count as one: evolutionary or genetic individuals. Many plants, and some animals, expand asexually. Some, like strawberries, send out runners, as strawberries do; some, like dandelions, generate asexually produced seed-like things. Now, note that I did not say that many plants *reproduce* asexually. I said that many plants *expand* asexually. That was intended to be neutral between treating it as reproduction and treating it as *growth*. That's because there is a genuine question about what we should consider to be an individual organism: should we consider the folk plants individual organisms, or should we take them to be mere parts of a much larger individual, which is marked off genetically? To make the question concrete: how many dandelions are in my yard? Many, as my neighbors surely think? Or just one?

In a wonderful short paper called "What are Dandelions and Aphids?" (1977) the biologist Daniel H. Janzen argues for the latter answer. He calls such an organism an 'evolutionary individual'. On the Janzen line, asexual expansion is growth; it is to get bigger, not to have babies. There is just one perennial dandelion in my yard, spatially scattered and many, many years old. It cannot be destroyed by a squirt of Roundup®. Quite generally, evolutionary individuals have different spatial boundaries and persistence conditions than the folk organisms that are their parts. As Janzen says of the EI aphid, "it is part of its growth pattern to spread itself very thinly over the surface of plants in the habitat, so thinly that a potential predator is unlikely to find much of it at once. Once much growth of the EI aphid has occurred, it is essentially indestructible as long as there is some food in the habitat" (1977, 588).

Why say this? Primarily because thinking of organisms this way enables importantly different insights into evaluations of fitness, success, competition, and so forth. Again, the details do not matter, and whether Janzen is right doesn't really matter either. What does matter are the following two points. First, there is a real question here: which is the better way to think of such organisms? Second, that question is one for biologists, not for philosophers, and not for the 'folk'.

The lesson from these cases should be clear. The Canberra Plan tells us to look around the world to see what, if anything, answers to our concepts. But we *also* need to look around in the world to see what concepts are useful in understanding it. We could sit around and do conceptual analysis on our ordinary, folk notion of

an organism. But what would be the point? What the right notion of an individual organism is—or, more likely, what the right notions *are*—depends upon the theoretical and explanatory concerns of the biologist. In short: conceptual analysis itself doesn't provide any guidance about what concepts we ought to be deploying. It says nothing about when—or how—we should revise or outright eliminate certain ways of thinking about the world. This is a significant limitation.

4. Limitation 2: Conceptual analysis by itself doesn't tell you whether anything in the world answers to the concept.

The second limitation, again, is that conceptual analysis alone doesn't tell you whether or not the concept is empty. This ought not be in any way surprising. After all, the whole point of the Canberra plan is to explicitly add a 'look around the world' step. Conceptual analysis alone obviously only tells us about our concepts, not whether anything answers to them. Nonetheless, it's worth explicitly emphasizing it, because philosophy's reliance on conceptual analysis is so entrenched that sometimes we almost forget that we're doing it. To illustrate this point, let's look at the literature on personal identity.

The question of personal identity is the question of what makes for the continued existence of a person over time. What kinds of changes can a person survive? In virtue of what am I the same person as that cranky toddler 40ish years ago?

There are of course a variety of views on this issue; here are simple versions of two prominent ones.

Bodily continuity: x at time $t1$ is the same person as y at time $t2$ just in case x and y have the same body.

Psychological continuity: x at time $t1$ is the same person as y at time $t2$ just in case x and y are psychologically continuous in the right sorts of ways.

According to the bodily continuity view, a person is an organism, a living body. So for a person to continue to exist through time and change is for that living body to continue to exist. According to the psychological continuity view, what's more important is the inside stuff, as it were: memories, continued experiences, continued plans

and projects, and so forth. Now, it's clear that satisfactorily fleshing out either view would require spelling out the relevant kinds of continuity in more detail. What exactly does it take for a *body* to persist over time? Bodies can survive some degree of physical change—we survive haircuts, surgeries, even amputations—but not just any physical change. Similarly, what exactly does 'psychological continuity' amount to? Since Locke, memory has been taken to be central; since Reid, it's been known that care is required with such a claim. So there is more work to be done to formulate a properly nuanced version of either account, but what is on the table is enough for my purposes.

Each of these views faces various objections, but in each case there is one that stands out. The standard objection to the bodily continuity view—one that is used to motivate the psychological continuity view—is the intuition that a person could survive getting a new body. There are a host of examples from fiction: Kafka's *Metamorphosis*, *Freaky Friday*, the *Star Trek* transporter.⁴ In such cases, we don't recoil as though the book or movie has asked us to imagine a round square, but rather simply enjoy the show. The thought is that this demonstrates that what matters to a person's continued existence is not particular physical underpinnings, but rather the kinds of psychological things mentioned above.

But, of course, what such thought experiments really show is that we *think* of ourselves as things that could in principle be transported into a new body. They are pieces of conceptual analysis, and as such only show something about our concept of a person. They don't show that the world in fact contains anything like that.

Let me press further with this point by considering a further move made by Richard Swinburne (1984). He conjoins the body-swapping thought experiments—the standard objection to bodily continuity views—with the standard objection to psychological continuity views.⁵ This is that such views apparently allow for the

⁴The thought is that the transporter does not move people's bodies from the ship to the planet's surface, but rather 'scans' them and recreates a new, qualitatively indiscernible, body from materials already on the planet's surface.

⁵ Again, these are standard objections; good places to read further about such matters include Perry 1978 and Shoemaker and

possibility of duplication. If, as the above line of thought suggests, the features central to the continued existence of a particular person can be ‘moved’ to a new body, then surely that could happen more than once. Consider the *Star Trek* transporter again. The conceit is that it records all of the data of the traveler, like uploading a CD and then burning a new copy. But of course once you have uploaded the CD, you can burn as many copies as you like. Similarly, if the transporter can create a version of Kirk on the planet’s surface, it can create more than one. (And, unsurprisingly, both the original series and *The Next Generation* exploited this for plot purposes.⁶) But this kind of duplication looks quite problematic, for familiar reasons that I will relegate to a footnote.⁷

Now, there are various responses available to both standard objections available. But we need not lose ourselves in the literature; all that matters for my purposes is that Swinburne takes them both to be successful, and concludes that persons can get new bodies but can’t be duplicated. On this basis, he endorses a soul theory of personal identity: we are immaterial, indivisible entities whose persistence over time cannot be further analyzed.

Swinburne 1984. A classic discussion of duplication, and an interesting solution, is Parfit 1971.

⁶ In the original series episode “The Enemy Within,” Kirk is duplicated but neither is exactly ‘right’ and the two are somehow ‘reintegrated’. In *The Next Generation* episode “Second Chances,” Riker is duplicated and the copy found many years later. Source: the interwebs.

⁷ Call the person who has yet to enter the machine ‘Kirk₁’, and call the two people who result ‘Kirk₂’ and ‘Kirk₃’. There are only three options: either Kirk₁ is identical to *both* Kirk₂ and Kirk₃; or Kirk₁ is identical to *one* of Kirk₂ and Kirk₃, but not both; or Kirk₁ is identical to neither. But on the face of it, none of these options is satisfactory. To say that Kirk₁ is identical to neither is to say that Kirk₁ dies upon entering the machine, but the defender of psychological continuity thinks that it is a legitimate method of transportation in non-duplication cases. To say that Kirk₁ survives as one but not the other would be arbitrary. And to say that Kirk₁ survives as both appears to lead to a violation of the transitivity of identity. If Kirk₁=Kirk₂ *and* Kirk₁= Kirk₃, then it must be that Kirk₂=Kirk₃... but of course Kirk₂≠Kirk₃. They are in different spatial locations and could have a conversation with each other!

And finally we get to the point. Swinburne is drawing a conclusion about the world based on arguments that, at their core, are pieces of conceptual analysis. They involve looking at our reactions to admittedly fantastical cases to draw conclusions about our concept of a person. Our *concept* of a person, not what persons really are. To make this concrete: I actually find Swinburne's view rather plausible *as a piece of conceptual analysis*. He may well be right that we think of ourselves as immaterial, indivisible entities. But I do not think that that is what we really are. Again, we need to recall *both* steps of the Canberra plan; conceptual analysis alone is not sufficient to tell us anything at all about what exists. Conceptual analysis is a useful tool for limning the boundaries of our concepts, not for telling us what answers to them. That is the second limitation.

And, as it happens, the first limitation of conceptual analysis rears its head again in this particular case. To see how, suppose for the sake of discussion both that Swinburne is right about our concept of a person, and wrong about the existence of immaterial souls. What follows? Well, by the letter of the Canberra plan, it follows that the concept is empty. There are no people.

But that is a bit extreme. It would seem, at least, that I am a person and so are you. So presumably the right conclusion would instead be that we ought to revise our concept of a person: we ought to think about what we ought to think a person is. That is, if matters are as I am supposing, the question for those investigating personal identity becomes: what would be a *useful* concept, close enough to our pre-theoretic one to count as a concept of a person rather than, say, a toaster, but not empty, and not so beholden to folk intuition? And, as we have learned, that is just not a question that conceptual analysis itself can answer. That was the first limitation: conceptual analysis is not a tool for choosing among available concepts. It says nothing about either when to revise or how to revise.

5. Conclusion

Conceptual analysis is a handy philosophical tool. It helps define the questions we are asking, and keep our philosophical theories connected to ordinary usage. But it needs to be used with care and attention, and it must not be taken to constitute a complete philosophical methodology. Conceptual analysis cannot tell us anything about what the world contains (second limitation), and while the

Canberra plan can, it cannot tell us anything about when or how to revise our existing concepts (first limitation).

Works Cited

- Gettier, Edmund. 1963. "Is justified true belief knowledge?" *Analysis* 23: 121-123.
- Jackson, Frank. 1998. *From Metaphysics to Ethics: A Defence of Conceptual Analysis*. Oxford: Oxford University Press.
- Janzen, Daniel H. 1977. "What are dandelions and aphids?" *The American Naturalist* 111: 586-589.
- Johnston, M. and Leslie, S.-J. 2012. "Concepts, Analysis, Generics and the Canberra Plan." *Philosophical Perspectives*, 26: 113–171.
- Lewis, David. 1981. "Are We Free to Break the Laws?" *Theoria* 47: 113-121.
- Lewis, David. 1972. "Psychophysical and Theoretical Identifications." *Australasian Journal of Philosophy* 50: 249-258.
- Parfit, Derek. 1971. "Personal Identity." *The Philosophical Review* 80: 3-27.
- Perry, John. 1978. *A Dialogue on Personal Identity and Immortality*. Indianapolis: Hackett.
- Schroeter, Laura. 2004. "The Limits of Conceptual Analysis." *Pacific Philosophical Quarterly* 85: 425-453.
- Shoemaker, Sydney, and Swinburne, Richard. 1984. *Personal Identity: Great Debates in Philosophy*. Oxford: Blackwell.
- Stace, Walter. 1952. *Religion and the Modern Mind*. Philadelphia: Lippincott.
- Strawson, Peter. 1959. *Individuals*. NY: Routledge.
- Swinburne, Richard. 1984. "Personal Identity: The Dualist Theory." In Swinburne and Shoemaker, eds., *Personal Identity*, 1-66.