What Are They Thinking?

Conversations with Australian Philosophers

Daniel Nellor

ASP

First published 2023 by
Australian Scholarly Publishing Pty Ltd
7 Lt Lothian St North
North Melbourne, Victoria 3051

tel: 61 3 93296963 / fax: 61 3 93295452
enquiry@scholarly.info / www.scholarly.com

ISBN: 978-1-922952-74-5

Cover design: Amelia Walker

Contents

Looking at Logic(s)

Greg Restall

Greg Restall is Shelby Cullom Davis Professor of Philosophy at the University of St Andrews in the United Kingdom. Before his arrival at the University of St Andrews in 2021, he held positions at the Australian National University, Macquarie University, and the University of Melbourne, where he was Professor of Philosophy since 2013. His research focuses on formal logic, the philosophy of logic, metaphysics, and philosophy of language, and even some philosophy of religion.

Greg Restall, what is philosophy?

In one sense it's indefinable, but a guiding idea that helps me think about what makes something philosophically interesting was put forward by the American philosopher Wilfrid Sellars, who said that 'the aim of philosophy is to understand how things, in the broadest possible sense of the term, hang together, in the broadest possible sense of the term.'

I like that.

I came to philosophy from the outside—my first training was in mathematics—and it was when I started asking those fundamental questions—'how does this work? How does it all hang together?'—that I started reading philosophy.

Is everyone a philosopher, really?

Well, you certainly find yourself doing philosophy whether or not you're stuck in an academic department. As soon as you ask the question 'does this fit with that?'—whether it's about some moral claim and how it relates

to a fact about the world, or some abstract piece of knowledge and how it connects with the ways we live—these are philosophical questions.

And within philosophy, what is logic?

Logic is also about how things hold together—so perhaps I've picked the definition of philosophy most amenable to a logician! In particular, logic is about the ways we represent the world to ourselves and to each other, in language and in thought and in other ways. But it's not about *what* we represent; rather, it's about how the thoughts that we have and the claims that we make hang together.

How they relate to one another.

Yes. Logic by itself won't tell you which claims are true, but it might be able to tell you that if *these* claims are true then these other claims must also be true; or that *this* is good evidence for *that*; or that this claim is inconsistent with this other claim and therefore you need to pick one—things like that.

Is logic kind of like the rules of the game for philosophy?

It's sometimes presented as a set of remedial skills you'll need if you're going to do philosophy. It's partly that. It's certainly useful to know how to spot a bad argument, because philosophers like to argue. It's good to know, for example, that 'if it's Friday, it's a weekday' does not mean 'if it's a weekday, it must be Friday'. These basic rules about what's a good move and what's a bad move are important. But the philosophical question is to ask: for those moves, how did they come to be good or bad? Is there one best set of rules? And what are we doing with those rules, anyway?

If I pick up a book on philosophical logic I'm likely to see a lot of symbols, and it might remind me of algebra. What is the relationship between logic and mathematics?

It's a good question, and it's one where philosophers take different views. One thing that relates logic to mathematics is that it's about finding patterns and expressing them as rules. Anything with rules and patterns can become mathematical in some sense: the rules of cricket, for example,

involve abstract structural features, as does music. A musical score is a kind of algebra: it picks out a pattern by abstracting certain features and ignoring others, which is why the same score can be given many different interpretations. Similarly with our thinking: logic is a way of attending to certain repeated features in our language or in our thoughts.

It seems to me that when we do mathematics—when we say, for example, that one plus one is two—there are logical questions and relations at play.

Yes, there are—for example, what do you mean by 'is'? Are you saying that one plus one is literally the same thing as two; or that 'one plus one' and 'two' are equal in quantity—and what do you mean by 'one' and 'two' anyway? These seem to function like names, but are they names for things? And if they are, what kind of things are they? We don't run into anything in the world called 'one'. It seems that we run into things like, say, oranges, and then we say things like: 'I've got an orange; I've got another orange; now I've got two oranges.' We put labels like 'one' and 'two' on these ways we have of engaging with the world, and then treat those labels as further things we can talk about. As a logician I'm interested in how this kind of thing arises. How do we go from just being in the world to saying things like 'the sun is shining' or 'the sun is not shining', and realising that propositions like these are things we can agree or disagree on, and that they can't both be true?

This is a bit of a funny question, but why should we be logical?

Well, it depends. What do you mean by logical?

Good question!

The reason I ask that is we often come with a bunch of connotations about what being logical means—for example, that being logical means I will only believe the things that logic drives me to, which is ridiculous. That's not how I come to ideas. Or we might use 'being logical' as a synonym for 'making sense' or 'not contradicting yourself' or for following a sort of

linear or calculating process—there's nothing wrong with any of that, but the way I approach the question 'why be logical' is to step back and look at how our ideas about logic connect with our practice. For me, logic is primarily about taking care with what we are saying.

And why should we take care with what we are saying?

Because we care about whether the things we're saying are true or false, and because we care about whether we have good reasons to believe the things we believe. We also care about whether we have good reasons to believe what other people are telling us.

Is there an ethical obligation to be answerable to the demands of logic?

There's certainly a *social* dimension to logic, because it involves language and communication. I don't get my concepts from myself, I get them from my community. You're in the same world as me, and therefore your ideas matter to me—partly because *you* matter to me, but also because sometimes we have to act together in the world, and you might have an opinion about something I care about. If your opinion is genuinely going to inform me about the world—that is, if it's not just an expression of fashion or advertising or something like that—then I need to be open to the possibility that you have reasons for your beliefs that I haven't thought of; and then I need to be able to tell what counts as a good reason and what counts as a bad reason. The tools of logic can help us to do that. So it's a fundamentally ethical matter in that sense.

What is the relationship between logic and reality? That is, does logic track reality—does it arise from how the world is—or is it something that we humans, with our rational capacities, impose upon *the world?*

This is an age-old distinction. As I get older—and also because I am by nature an agreeable fellow—I like to agree with both sides, which is very annoying for people who want to argue with me! I want to say that on the one hand, the reason there is something inconsistent about saying 'it is sunny outside' at the same time as saying 'it is not sunny outside' is that

it *can't* be both sunny and not sunny at the same time, and this is simply a fact about the world. But on the other hand, what happens if we zoom into some of the concepts we're using here: like the word 'not', for example? What makes the word 'not' do the job it is doing when I say 'it is *not* sunny outside'? One way of looking at this is to say that the word 'not' is designed to express disagreement, and it exists because if it didn't, we would all go around agreeing with each other all the time, or at least not noticing that we were disagreeing and having no way to bring disagreements to the surface. The concept 'not' can do this job, if we play the 'not' game in a particular way. So, if we look at the question this way, then it's easy to see logical operators like 'not' as something we impose upon the world.

So, you're suggesting that logical principles such as the law of non-contradiction—which is the principle that a thing can't be true and not true at the same time—might be there in our logic because they help us to get by in the world, in this case by giving us a way to express disagreement.

Yes.

But isn't it also the case that we need a way to express disagreement because of the way the world is? That is, I need to be able to deny the sun is shining because sometimes, in the real world, it just isn't!

Yes, and that's the other side of the age-old distinction. I've just suggested that we humans wanted to be able to do certain things in the world and so we designed our logic and language practices accordingly. We might even, in principle, have designed them differently. But the response to this is to say hang on: surely we designed our practice in this way because it puts us in touch with the world as it is. And surely now that we have these concepts of contradictoriness or inconsistency or negation or exclusion, we can use them to look at the world and notice its real features, such as that it's being sunny excludes it's not being sunny. Again, I agree with both sides of this argument! I do agree that there seems to be something universal about logic. If aliens arrived and could communicate with us, I wouldn't be

surprised if they thought, like we do, that two plus two equals four, or that if p is true then the negation of p isn't true, or whatever, because that's just how the universe works. But what I'm really interested in is what the rules of the game are, how we came to have these rules, and most importantly, how having logical rules is something that we, as humans, can do.

We've discussed the relationship between logic and the community. What about diverse communities? Does cultural diversity bring about logical diversity?

If we look at, say, Buddhist logic, or Indian or Arabic or Chinese logic, it's true that we sometimes find different traditions analysing things in different ways. In the end, all of these logical theorists are theorising about our everyday practices, and our everyday practices are messy. So, I think it's possible for the experiences of different communities to be theorised differently, and there may be no fact of the matter about which of these theories is 'better'. But of course, this is in tension with the thought that logic deals with claims about the world, and if something counts for me as a claim about the world then it has to be a claim about *my* world, even if it's about a part of the world that I haven't experienced. So, I don't want to say that we live in completely different, hermetically sealed conceptual worlds.

Because then we wouldn't recognise each other as making claims about the world at all.

That's right. And of course, plenty of bi-cultural people live in more than one cultural world at the same time and find ways to fit them together. But it's certainly true that if I'm dealing with the logical theory of another culture, I need to be very patient and not assume that my first stab at understanding it is correct.

Sometimes the philosophical ideas that emerge from different cultures include different metaphysical pictures of the world: for example, the Buddhist concept of the 'groundlessness of all being' holds that everything is connected, everything is dependent on something else, and

there's no ultimate, lasting reality that is not subject to change. Does the background metaphysics of a particular culture inform its logic?

Some Buddhists are more metaphysical than others when they talk about the relationship between the groundlessness of all being and the contingency of all claims. Similarly, non-Buddhist logicians disagree with each other as to the metaphysical status of *their* logical claims. But I'm hesitant about the suggestion that different metaphysics give you different logics, because part of the self-conception of logic is that its rules and structures can be applied to debates about metaphysics.

Which means that the rules and structures of logic need to transcend the various metaphysical pictures of the world.

Logic at least claims to do that—whether it's successful is another question. It might be suggested that every logic carries an implicit metaphysics; but one way you could respond to that challenge is to say that if two people have very different metaphysical conceptions of the world, and if they disagree with one another, then they at least have a starting point from which to talk about what counts as an agreement or a disagreement and so on.

And then they're doing logic.

That's right.

You're currently working on a big project, with scholars from many different disciplines, on the idea of 'social hierarchies'. I'm interested in how you as a logician fit into this. What is the relationship between logic and power? Is the way we understand logic tied up with our power relations? Or can good logical practice help to undermine unjust power relations?

Again, both of those things are true. As a logician, I'm interested in social hierarchies because my work on logic and language is deeply connected to our practices as human beings in society. I'm interested in how the moves that are available to us in dialogue and discussion are affected by our position in social space: whether I'm an expert or a student, for example, or a male or

a female, or whether I'm from a dominant culture or an oppressed one—all of these social positions we find ourselves in affect how we communicate, and how we interpret one another.

And how might thinking about this inform the way we think about logic?

The stuff logicians tend to think about can seem pretty dry and abstract, and in one sense, value-free. Logicians think about how words like 'and', 'if' and 'not' behave, for example. But I want to suggest that these 'primitive' concepts—these basic building blocks of thought—arise out of a practice which is not really value-free at all. I'm interested in the idea that there are other concepts which are equally or even more primitive than the ones logicians typically focus on. For example, I teach my first year logic students about quantifiers: words like 'all' and 'some'; and when I do that we might discuss the obvious difference between a sentence like 'all birds lay eggs' and a sentence like 'some birds lay eggs'. But think about the content of that sentence. The idea that 'birds lay eggs' is something people learn way before they come to first year logic—we learn it in kindergarten, or earlier. It's a generic, like 'frogs are green' or 'pigs like mud'. But there are other generics also worth thinking about: 'Men are violent.' 'Women are submissive.' 'Cows are food.' Generics are default generalisations or expectations, and some of them are really weird. We say 'birds lay eggs' but we don't say 'birds are female'. But half of all birds are female and not all of them lay eggs.

Where do you think these generics come from?

I think they're part of our explanatory practice, our need to make sense of the world. We say, for example, 'mosquitos carry the West Nile virus'. But what do we mean by this? We don't mean that *all* mosquitos carry the West Nile virus, or even that most do—but if I come down with the West Nile virus and a doctor tells me I was bitten by a mosquito, then I'm going to say, well, that explains it. I don't need any further explanation.

So a generic is a kind of shorthand that helps us understand the world we live in.

Yes.

But we can see how this can go wrong. When I studied logic I learned that if you get a quantifier like 'all' or 'some' wrong somewhere in the chain, it can have huge ramifications down the track. How does this connect to questions about 'social hierarchies'?

If I say 'black lives matter', what am I saying? Someone might reply, 'no, *all* lives matter.'

Is this just a confusion of quantifiers? After all, no-one who says 'black lives matter' really means 'only *black lives matter'.*

It's an example of how different kinds of communication can be going on when we use generic expressions. The same person who objects to my saying 'black lives matter' might also object to a comment like 'men are dangerous'. He might say: 'I'm a man, and I'm not dangerous.' In this case he's assuming that my use of the generic 'men are dangerous' has a kind of universal power. Generics play an important communicative role—they help us describe the world and inform one another about things—but they also seem to have a coordinating function. If I say 'birds lay eggs' my claim is not really about what proportion of birds lay eggs; it's more about how we should expect things to happen. Generics seem to coordinate default expectations and help shape the frames of reference through which we understand things.

For better or worse.

Exactly. One of the reasons generics can be so hard to get agreement on is because they're often expressing very different orientations to the world.

In general, to what extent can fixing logical errors in our thinking and practice improve ethical outcomes?

There's a sort of naive view of what it is to be logical that says if you just put all the facts into the logic machine you can grind out a result and it will always be right. This is actually a dangerous ceding of responsibility. We teach our first year logic students that if you have a valid argument—that

is, an argument where if the premises are true then the conclusion must also be true—it's always an option to look at the conclusion and say 'no—that's rubbish!'

And then you need to look back at the premises.

That's right. You might immediately see that you have a good reason not to accept a certain premise. Or you might identify the problems you have with the conclusion and use this to trace back your reasoning and see something you couldn't see before, that one or other of the premises is now questionable. Or you might realise that your 'logic machine' is broken. The reason we study logic is not so that we can feed all of our reasoning capacity into the machine and pump out the right answers, but rather that we can become good mechanics, able to lift up the hood and see what's going on. And yes, doing this well or badly can have ethical consequences.

Let's move on to a particular view of logic that you have argued for, and that's logical pluralism. In a paper you co-wrote with J.C. Beall, you start by saying that over the last eighty years or so the subject matter of philosophical logic, or at least its emphasis, has changed from the search for logical truth to an exploration of logical consequence. Could you explain what you mean by this?

I said earlier that for me logic and philosophy are about how things hang together: does this contradict that? Does this lead to that? These are questions about logical consequence: that is, what follows from what, and how? Most academic logic these days is about trying to figure out how logical consequence works. Logic used to be about discovering eternal laws, the truth of which could be discovered through logic alone; but today the emphasis is more on how claims about the world relate to one another.

The term 'logical pluralism' suggests there are several different logics. Would you agree that most people's intuition is that there is just one true logic?

I think that's the default position, in the same way that most people think there's one true arithmetic.

And the task of the logician is to discover the one true logic.

Yes, and expound on it.

But logical pluralism denies this.

Right.

When we talk about different 'logics', we often define them in opposition to 'classical logic'. Can you say a bit about what this is?

Classical logic involves taking the individual claims we make about the world—propositions—and looking at how we combine them using words like 'and', 'if', 'not', and so on. It's a theory about how these fundamental concepts work, and it starts by saying that when we make a claim about the world the claim is either true or false.

It's either sunny or it isn't.

Right. Now let's assume I have an argument consisting of three claims, each of which is either true or false. I might look at this argument and say: well, I can see that if the first claim is true then the second must also be true, and if the second claim is true then the third must be false. That means I can also see that the first and the third can't both be true.

Because if the first is true that makes the second true, which makes the third false.

Correct. Now, according to classical logic, we can take an argument like this and list out all the possibilities—true-true-false, false-false-true, false-true-false, and so on—and classical logic says that if there is no way to make the premises true and the conclusion false, then the argument is valid.

Which doesn't mean, of course, that the conclusion is actually true; 'valid' just means that if the premises are true, then the conclusion must also be true.

That's right, it's about the structure of the argument, not the truth or otherwise of the conclusion.

So to take a simple example: if I say 'All humans are mortal; Samantha is a human; therefore Samantha is mortal', it's clear that if the first two premises are true, then the conclusion must also be true. The combination 'true-true-false' is not possible for this argument; therefore it's valid.

For classical logic, that's right.

And it works because every claim must be either true or false.

Yes, classical logic says that every claim is either true or its negation is true. It's either sunny, or it's not sunny. Either *p*, or not-*p*. But there have always been people who said wait—I can think of plenty of claims which, all things considered, I don't think are true, but I don't think their negations are true either.

And as soon as you've said that you've left classical logic behind.

Yes, because then you're saying that not everything needs to be either true or false. A non-classical logician might have a different story to tell about the truth values of propositions, and about many other things as well.

But to be a non-classical logician is not the same as being a logical pluralist, is it?

No, because you might think that although classical logic is wrong, there is still only one true logic. And you probably think that your particular non-classical logic is the right one!

So how is a logical pluralist like yourself different?

I think that the verdicts rendered by two different logics can *both* be correct. Let's focus on a concrete example. Suppose we have the following argument: Premise one—it's raining. Premise two—it's not raining. Conclusion—therefore, it's Saturday. Now, that's obviously a crazy argument. It's an argument from contradictory premises—it's raining and it's not raining—

to a conclusion which has nothing whatsoever to do with the premises. But if you put it into the classical logic machine it comes out valid.

Why?

Remember that in classical logic, an argument is valid if there's never a case where the premises are true and the conclusion false. For the crazy argument I just outlined, there's never a case where the premises are true and the conclusion false, because there's never a case where the premises are true. It wins by default.

I can see why some people might have a problem with classical logic!

Exactly. Ever since Aristotle there have been people who said that there should be a constraint in logic that says that the conclusion should be about what the premises are about. For an argument to be valid, the conclusion should be somehow *relevant* to its premises. And so, there was a tremendous argy bargy for many years between relevant logicians and classical logicians. But what makes me a pluralist is that I went back to the beginning and said hang on—what do we mean by 'validity' anyway?

And what do you think we mean by 'validity'?

I think we mean something like 'good'. A valid argument is an argument we consider to be good for some reason. And aren't there many different ways that things can be good? Can't an argument be good in one sense and not good in another? If a classical logician says to me, 'look, here we have an argument where it's not possible to make the premises true and the conclusion false'; I might say, yes, that tells me something important about the argument, so I'm happy to call it good. On the other hand, I might notice that some arguments which are good in that classical way also have a failing that's worth keeping track of: that their conclusions have nothing to do with their premises. And so I might think that there might also be another, different way for arguments to be good: that they be 'subject-preserving' or 'relevant'. And there might be many other ways for arguments to be good, too. So, I think that any particular argument can be

seen to have *this* virtue but not *that* one, and therefore be good in *this* way but not *that*. This makes me a pluralist: I don't think there's an ultimate answer as to which is correct.

But if an argument is good in a certain sense and bad in another sense, how do we know, when we infer a conclusion from a set of premises, whether we've made a mistake or not?

It depends what kind of mistake you're talking about. The notion of a 'mistake'—like the notion of 'good'—is normative. It's about how we value things. When you ask whether someone has made a mistake, one thing you might be asking is whether or not their conclusion is true. But that's not a question for logic.

Right, because logic is about how we get from the premises to the conclusion.

That's right. You might ask 'well, is this particular argument a good way of getting from these premises to that conclusion?' But again, I'll want you to tell me what you mean by 'good'. Good for what? Good in what sense? As a logical pluralist, I'm asking people to pay more attention to the different kinds of evaluative notions we use.

You're clear in your paper that being a logical pluralist is not about saying 'anything goes'.

No. A logic offers a set of criteria for saying that certain arguments are good and certain arguments are not good. What we pluralists care about is whether the notion of 'good' that is being employed in a particular logic is important and worth keeping track of.

Okay, but if it's not the case that anything goes then doesn't that mean there must be something that makes this *logic acceptable and this other logic not acceptable? Isn't there some criteria which transcends all logics and rules some in and some out? And if so, is that not the 'one true logic'?*

Whatever 'that' is, it's not a logic. Maybe it's a set of principles for discerning

what kinds of things can count as virtues for arguments—a set of claims *about* logic, but not a logic itself. A logic is a theory of valid consequence. A pluralist simply wants to say there's more than one of those, and we don't have to make some ultimate choice between them.

But is it possible that as a pluralist you could endorse two logics, one of which validates an argument to the conclusion A, and the other of which validates an argument to the conclusion not-A?

No, I don't think you should endorse two logics that lead you from premises you accept to contradictory conclusions. That might be a kind of constraint on which logics are acceptable to a pluralist.

So just as reaching a crazy conclusion might lead a classical logician to re-examine her premises, the realisation that two different logics can validly lead from acceptable premises to contradictory conclusions might lead the pluralist to think again.

Yes, it's exactly the same phenomenon at a higher level. But it's important to remember that there are logicians, such as my friend Graham Priest, who think that some contradictions are true. I don't myself think that any contradictions are true, so when two logics lead me to contradictory conclusions I'm going to see them as giving me conflicting advice, and that's a problem. But Graham, who thinks that sometimes A and not-A can both be true, might be fine with that.

Are more logicians pluralists, or non-pluralists?

I don't know the exact numbers! But I certainly think there are more pluralists than there used to be. And people can be pluralist in different ways.

You co-wrote another paper, this time with Francesco Berto, called 'Negation on the Australian Plan'. This is about how negation—the word 'not'—should be understood as working in logic. In this paper you're opposing an 'American Plan' for negation. What's motivating this debate?

Earlier we were talking about classical logic, which says that an argument is valid if the premises can never be true while the conclusion is false. We discussed a problem with this, which is that contradictory premises can lead to any conclusion whatsoever, and the argument can still be valid.

So to recap: if the first premise says 'p' and the second premise says 'not-p', then the conclusion can say anything it likes—'the moon is made of green cheese'—and the argument turns out to be valid because even though the conclusion is clearly false, the two premises can never be true.

Correct, which is why the relevantists say no, that can't be right—the conclusion of a valid argument should be in some way relevant to the premises. Now, some people from the US think that the best way to avoid these problems is to expand the number of truth values we give to propositions.

The truth value of a proposition being whether it's true or false.

That's right. Traditionally truth value has been two-fold—that is, any proposition must be either true or false. But the American Plan thinks of truth value as four-fold: a proposition can be either true, or false, or both true and false, or neither true nor false.

And how does this help with the relevance problem?

It affects the way negation works. If truth value is two-fold then the negation of any true proposition must be false. If *p* is true then not-*p* must be false. Conversely, if not-p is true then p must be false. So for classical logic, when it comes to that strange argument about the moon and green cheese, it can never be the case that both premises are true, because they contradict one another, so no matter how crazy the conclusion is, the argument remains valid.

Because it's not a case of the premises being true and the conclusion false.

Right. But what if the American Plan is right and premises can be both

true *and* false? If *p* is both true and false, then the negation of *p*—not-*p*—is also both true and false. This means that both of the premises in the crazy argument, if they're both true and false, are at least true. And then we *do* have a case of the premises being true and the conclusion false, and so the crazy argument is no longer valid.

I see. So the American Plan is one way of avoiding having to call nonsense arguments like the one about green cheese valid, by understanding negation as not always flipping the truth value of propositions from true to false or vice versa.

That's right, and it's simple and elegant—but it's not the Australian way.

So the Australian Plan has a different way of dealing with the relevance question?

Yes. On the Australian Plan, relevance is determined by situations or circumstances. In classical logic, every proposition in the entire world is either true or false. Faced with any claim whatsoever, we have two options: to call it true, or to call it false. On the American Plan there are, on top of these two options, two more options for every claim: it might be both true and false, or it might be neither true nor false. But the Australian Plan says something different. It says, let's not think about the entire world: let's just think about claims as being true or false only in much more limited circumstances.

Could you give an example?

Think about the circumstances of my study right now. The circumstances in my study make certain things true: that I'm wearing a red jumper, for example, and that I'm sitting down. These things are made true by the current situation in my study, in which I am indeed wearing a red jumper, and sitting down. But the situation in my study has no bearing on the question, for example, of whether the President of the United States is currently sleeping or not. What settles *that* is some other situation. In the bit of the world that is my study, it is not true that the President is sleeping,

or that he's not sleeping. What makes it true that he's sleeping, if he is, is some other bit of the world.

And how does this affect negation?

Classical logicians don't like talking about bits of worlds; they like talking about entire worlds. So in classical logic you can say things like 'if Greg is wearing a red jumper, then either the President is asleep or he is awake'. And this is true in classical logic. In fact, it can't be false, because the President *must* be either asleep or awake. But this invites the thought that my wearing a red jumper has something to do with the fact that the President is asleep or awake. But it doesn't. My wearing a red jumper does not make it true that the President is asleep, if he is, and does not make it true that the President is awake, if he is. The situation in my study simply isn't big enough to make either of those claims true or false. But the situation in my study *is* big enough to make it true that I am not standing up right now. Why? In our paper we suggest that the best way to understand this is in terms of compatibility. Any possible situation which involves me standing is incompatible with the situation in my study right now.

Because you're sitting.

That's right.

Whereas the situation in your study is entirely compatible with the President of the United States being either asleep or awake.

Or indeed with him being both asleep and awake at the same time.

Which is impossible, but it's not made possible or impossible by the situation in your study.

That's right. What makes either of those options true has nothing to do with my study. What makes my current sitting or standing true has everything to do with the bit of the world that is my study, because that's where I am.

So on the Australian Plan, a claim of negation is actually a claim of incompatibility. When I say 'it's not sunny', I'm claiming that the

situation in the bit of the world to which I'm referring is incompatible with its being sunny.

That's right. It's a way of thinking about negation that is designed to pay attention to subject matter—to what claims are about. It's a way of requiring conclusions to be relevant to premises. It's trying to capture an intuition about the crazy green cheese argument: when you have a premise, *p*, and then the negation of that premise, not-*p*, the two situations in which each of those premises is true are in some sense incompatible with each other.

But each of those situations is compatible with the moon being made of green cheese, or not.

Right, which means the premises in the green cheese argument are not saying anything about the conclusion; they don't make true anything about the moon, and so we avoid having to call the argument valid.

On the Australian Plan, do you preserve the two-fold truth value?

Yes. The fact that the situation in my study does not make it *true* that the President is sleeping, does not mean that in this situation the proposition 'the President is sleeping' is *false*. And this is not because it is both true and false, or neither true nor false, as the American Plan would have it: no, it is simply that the situation in my study is incomplete, and doesn't have an opinion on the question.

In the world of logic, where does the Australian Plan sit?

Oh, it's a very niche subject. But those of us who really care know we're on to the truth!

Paying Attention

Christopher Cordner

Christopher Cordner was Associate Professor of Philosophy at the University of Melbourne until 2021. He is now Honorary Principal Fellow in the School of Historical and Philosophical Studies at the University. His main area of interest is ethics, including its classical Greek and Christian tributaries, and its overlap with art and aesthetics.

Christopher Cordner, what is philosophy?

The best way to give someone an answer to that question is to engage them in some actual philosophy! One reason for that (there are others) is that 'what is philosophy?' is itself a philosophical question—and so is immediately 'up' for philosophical contestation. Even so, I'm inclined to respond to your question by saying: 'Philosophy attempts to make sense of the world and of ourselves in it.' But there are problems with saying that. It is at best a very vague statement until you are into some philosophizing and begin to get some sense of what the answer might actually come to. And anyway, should we accept the apparent implication that until philosophy comes along 'the world and ourselves in it' *doesn't* make sense? And there's another problem: creative art of many kinds also undertakes to make some kind of sense of the world and of ourselves in it, and so do various religious outlooks, and so can scientific inquiries. But that doesn't mean that art *is* philosophy, or that a religious orientation or a scientific account of something is philosophy. And then how we are to demarcate what constitutes a philosophical making-sense from other ways of doing that is a tricky matter. But all *that* now said, I think the answer I mooted still does offer us *something*.