What is an introduction?

Metaphysics is one of the traditional four main branches of philosophy, alongside ethics, logic, and epistemology. It is an ancient subject but one that continues to arouse curiosity. It holds an attraction for many who have only a basic inkling of what it is but are keen to know more.

For some, it is associated with the mystical or religious. For others, it is known through the metaphysical poets who talk of love and spirituality. This book will aim to introduce the uninitiated to how metaphysics is understood and practised by philosophers. Many introductions to the topic begin with a consideration of what metaphysics is and how its truths can be known. But this itself is one of the most difficult and contentious questions, and the reader could quickly become bogged down and lose interest. This book is therefore written back to front. The question of what metaphysics is and how it is justified will be left to the very last. The best way to understand an activity is often through doing it rather than theorizing about it. In that case, we start by doing some metaphysics: considering some seemingly simple little questions but which concern the fundamental nature of reality.

We will go through a variety of issues with only a few technical concepts and terms. By the end, there should be a fair grasp of the problems around substance, properties, changes, causes,

possibilities, time, personal identity, nothingness, and emergence. It is hoped that the book will not intimidate its readers in a way that many philosophy books – particularly in metaphysics – can.

Often, the ideas, concepts, and questions of metaphysics sound easy – childish even. What are objects? Do colours and shapes have some form of existence? What is it for one thing to cause another rather than just being associated with it? What is possible? Does time pass? Do absences, holes, lackings, and nothingnesses have any form of positive existence at all? To some, these seem like silly questions, but for others they are at the core of what philosophy is all about. And those who see it that way often get a sense that the issues these questions raise are the most fundamental and profound about which humans can think. Metaphysics is the subject among all others that inspires the sense of wonder in us, and for that reason some think that doing metaphysics is the most valuable use we could make of our time.

If you have made it this far, perhaps metaphysics has already captured your imagination and your curiosity. In that case, we should begin forthwith on our little tour of the metaphysical furniture of the world. But where to begin? Philosophers never really know. The things they worry about are often interconnected. To understand one issue, you need first to understand another. Yet we have to say the same about the second issue as well: to understand it, you need to understand a third, and so on. And this seems to be true no matter where we start. Sometimes an understanding of the world comes only by grasping the whole, which makes it hard to explain the problems of philosophy in a neat sequence, as books must inevitably try to do. Where we begin is thus to an extent arbitrary.

Chapter 1
What is a table?

When I look at the world around me, I see that I am surrounded by all sorts of things. I see a table and two chairs, buildings, an aeroplane, a box of paper clips, pens, a dog, people, and a wide variety of other kinds of things. But this is a book about metaphysics, and in metaphysics we are concerned with the nature of things in very general terms. I am tempted to say, as a metaphysician, that all of these things I have listed are particular things, or groups or kinds of them. The notion of a particular is very important to us. I want to know that the pen on the table is my particular one rather than someone else's, or that the woman in the room really is my wife rather than her identical twin sister. To understand the importance of these issues, we need to probe them more deeply.

In front of me stands a table that I can see, feel, and hear if I rap my knuckles on it. I have no doubt that it – the table – exists. But now I will start the philosophical questions. What is this thing? What is the nature of its existence? Is the table something I know through experience or do my senses reveal to me something else? After all, when I look at it, I see its colour: the brownness of the wood. And when I feel it, I feel its hardness. Brownness, hardness, four-leggedness, and so on, are the qualities or properties of the table. One might then be tempted to say that I do not know the table itself but only its properties. Does that then mean that the table is

an underlying something about which I know nothing? Its properties seem wrapped around it and impossible to strip away.

What goes for tables, goes for other particular things too. There is nothing special in the choice of a table as my example. In the cases of coins, motor cars, books, cats, and trees, I know them only through knowing their qualities. I see their shape, their colour, I can feel their texture, smell their fragrance, and so on. The nature of these properties of things – redness, roundness, hardness, smelliness, and so on – will be the topic of the next chapter. But we really cannot avoid mentioning properties as soon as we mention the particulars to which they attach.

The more things change, the more they stay the same

Now why would I suggest that the table is something other than the brownness, hardness, and four-leggedness that I can see in front of me? One reason is that I could imagine these properties changing while the table remains the same particular that it was. I could paint the table white, for instance, because it fits in better with the decor of my office. If I did that, then it would still be one and the same table, it would simply have changed its appearance. Something will have changed, while something has remained the same.

In philosophy, we see that all sorts of confusion can reign if we speak loosely of it being the *same* table, so we employ an important distinction. We can say that something has changed *qualitatively* even though it has remained *numerically* the same. So the table can be different in its qualities – it was brown and now it is white – but it remains one and the same thing. The table that was brown is now the table that is white. Imagine if a visitor comes into my room and asks what's happened to my old brown table. It's perfectly acceptable for me to respond that it's still here: it's just that they didn't recognize it because I had painted it. Being

one and the same, despite such changes in qualities, is what we mean by *numerical* sameness (the topic of change will be explored more in Chapter 4).

It is this consideration that leads me to think that the table itself cannot be the same thing as its properties. At least some of them could change and yet it would still be the same table. So when I look at and feel the properties of the table, I am observing just that – its properties – and not the table itself. But what, then, is the table, if it is not its properties?

Here is a suggestion. The table is something that underlies the properties and holds them all together in one place. It is something I cannot see or touch, because all I experience is a thing's properties, but I know it is there through my rational thinking. When I move the table across the room, for instance, all of its properties move with it. They are clustered together in a semi-permanent way. It is not as if the brownness and hardness of the table can move but the four-leggedness can get left behind. I say that the properties are clustered only *semi*-permanently, though. As we have seen, some properties can be shed from the cluster and new ones take their place, so we cannot be absolutely strict and say that the properties are bound together inseparably. The brownness can be shed and replaced by whiteness.

Such a view of particulars may be best understood through the metaphor of a pin cushion that is used to hold pins together in one place. The pins represent the properties of an object and the cushion represents the particular itself. Some call this a *substratum* view of particulars, where the pin cushion is the substratum that underlies all the properties on view. One pin stands for the brownness of the table, another stands for its hardness, and a third stands for its weight, another its height, and so on for every single property the table has. And if we could strip these away – mentally, through a process of abstraction – we would come to understand that the thing itself is separate from

them and is that in which they all inhere. Of course, when you remove all the pins from a real pin cushion, you are still left with something that you can see and touch. But remember that our metaphorical pin cushion, when all its pins have been removed, is a particular that has been stripped of all its properties so that we can think of what the table itself is. And without properties, it couldn't therefore look or feel like anything.

Consider, for instance, a cat. We can think of it without its blackness; for that is a property and we want to know what the thing is that underlies all its properties. But removing its blackness isn't like skinning a cat. As well as removing its colour, we also have to take away its shape, as that is just another property like the rest, and so is its four-leggedness, smelliness, and furriness. Take all those away and we could well wonder what this underlying substratum really is. It would have to be invisible. It would have no length, breadth, or height, and no colour or solidity. There would be a bareness to it that may really make us start to wonder whether we have anything at all.

Philosophers are notorious for working out all the implications of an idea. But they don't necessarily always accept those implications. Sometimes a consequence is so ridiculous that it can be taken as good grounds for rejecting the initial supposition. Such a counterintuitive consequence will have reduced the supposition from which it sprang to absurdity. Perhaps we can say that's happened in this case. It was suggested that the particular had to be something other than its properties. But once we started to abstract away the properties of the cat from the cat itself, we realized that it would leave hardly anything. Our substratum-cat seems to be nothing at all. It has no weight, no colour, no extension in space, and so on. And this starts to look like a non-thing. Isn't it the case that everything that exists has properties? It is not as if 'bare' particulars could exist and that some of them were just fortunate enough to accidentally acquire properties. Certainly every physical thing that ever has and ever

Metaphysics

6

will exist has some shape or weight or feature. And to talk as if the thing can in some way exist independently of those properties was perhaps the mistake that led us to absurdity.

Bundles of properties

Let us, in that case, consider a different approach. If there can be no 'bare' particulars, existing without having properties, then we might want to think again of the cluster or bundle of properties with which we began. When in our minds we stripped away those properties, in a process of abstraction, the fear was that we were left with nothing at all. So shouldn't we then just countenance the possibility that there is nothing more to a particular than that bundle of properties? If there really is no remainder once all the properties have been removed, then we know that our particular cannot be more than them. The bundle view is that particulars can be accounted for in terms only of properties. How plausible is this view?

There are a couple of problems associated with it, which come from the problem of change that we already discussed. If a thing were just a collection of properties, it couldn't survive any change. If one property were lost and another gained, we would have a different collection: for I am assuming that what makes a collection the same thing at different times is that it is composed of the same component things. Consequently, two collections are different if the things collected within them are different. And clearly, the particulars that interest us change all the time while remaining (numerically) the same. A cat changes its shape frequently. Sometimes it is lying out flat, other times it is rolled up in a ball, and then it might be running around, changing its shape continuously. How can the cat be just a collection of properties when they change all the time?

It may be possible to answer this objection, though. Perhaps we should think of a thing as a series of bundles of properties, united

by a degree of continuity. So while the table can be changed and painted white, it keeps roughly the same weight, height, and physical position. I am assuming the physical position of an object is one of its properties, and clearly it is a pretty important one in this context. I am confident the white table is the same thing as the previous brown table in no small part because I find it in the same room. And if it has moved, I expect that it did so gradually by passing through a series of locations between where it started out and where it ended up. While the cat changes shape rapidly, it keeps the same colour, furriness, smell, and, importantly, it is in the same place; or if it has changed its position, it has done so through a series of locations. We could say, therefore, that while the bundles of properties come and go, a particular thing is a succession of such bundles with an appropriate continuity running throughout.

Metaphysics

There are a number of other difficulties to be faced, but before going on to consider one of them, it is worth mentioning what might be a big advantage of this bundle view. The first account we considered was one in which particulars were underlying substrata that held the properties of a thing together. To account for particular objects such as a table, a chair, a dog, and a tree, we had two kinds of ingredients. We had a thing's properties and its substratum. But with this new bundle theory, it seems that we need only one kind of thing. We just have the properties and, when they come in a bundle or a continuous sequence of such bundles, we say that we thereby have a particular object. So where we previously needed two elements, we now have only one. Another way of looking at this is to say that the notion of substratum has been reduced away entirely in other terms. Objects would just be nothing more than bundles of properties, appropriately arranged.

The second theory is thus a simpler one in so far as it invokes fewer kinds of entity. The unknowable formless substratum seemed to give us nothing extra: if the bundle theory is correct,

then the substratum is dispensable. Now there is no particular reason why a simpler and more economical theory is more likely to be true than a complex and uneconomical one, but philosophers prefer the simple ones. Certainly, there seems no reason to tolerate redundancy in one's theory of the world because any redundant elements are clearly not needed for the account to work. They serve no purpose.

Identical twins

The bundle theory looks simpler than a substratum view, therefore. But is it too simple? Would it have enough resources to deliver all we want of a particular thing? There is one consideration that suggests not. A particular, we are told by this theory, is just a collection of properties. A snooker ball, for instance, is just a bundle of the properties red, spherical, shiny, 52.5 millimetres in diameter, and so on. The problem for the theory, however, is that there could be another object with exactly those properties. Indeed, for the game of snooker to be fair, there should be many red balls with those same properties: they are standardized. The theory has a difficulty here, however. It tells us that a particular just is the bundle. But then, if we have the same bundle, it implies that we have the same object. In other words, there could not, on this theory, be more than one object that is the same bundle of properties.

It might be said that this objection is a mere technicality that doesn't really matter. Couldn't it just be that, as a matter of fact, two distinct objects never really do share all the same properties? Even tables that are mass-manufactured will have some very slight difference in weight, colour, or even just the pattern of fine, microscopic scratches on the surface. Our snooker balls need only be close enough in their properties for the game to be playable fairly so they too can have some slight differences. This response misses the point of a philosophical theory, however. This was supposed to be an account of what it is

What is a table?

9

to be a particular thing. The truth of that theory should not have to rely on luck working out for it, such that every particular thing just happens to be a different bundle. It does seem at least a possibility that two things could share all their properties. And if, as the theory states, particulars are only and nothing more than bundles of properties, then it is inconsistent with that possibility. Two particulars with the same properties collapse into one.

There are two possible ways out for the bundle theorist but both have problems. The first apparent solution is to say that there is a reason in principle why two particulars could not share all their properties. If one allows relational properties, then these arguably must differ because they allow spatiotemporal location to come into the equation. The following example illustrates what is meant by a relational property. Even if all the red snooker balls are indistinguishable when you inspect them, perhaps one is just 20 centimetres from the bottom-right pocket of the snooker table, while the other is 30 centimetres from it. One ball has the relational property of being 20 centimetres from the pocket, while the other has the relational property of being 30 centimetres from the same pocket. Assuming no two entirely distinct particulars can occupy the same space at the same time, then it seems that all things will bear a unique set of relational properties.

Here is the problem with this proposal. There is no guarantee that distinct things really will have different relational properties unless we reintroduce particulars into our metaphysics. This is why. Should we think of position in space (and time) as an absolute or relative matter? If it were absolute, it would suggest that there is some kind of particularity to spatial positions. A position would be a particular. The notion of a particular – one that is not defined as a bundle of qualities – will have come back into the theory. That's no good because we were looking to eliminate particulars in terms of bundles of properties.

10

So do we instead define spatial positions in relation to each other? The problem with doing so is that there is at least the possibility that the space of a universe has a line of symmetry; and thus places in corresponding positions either side of the line of symmetry would bear an identical set of relations to all the other places within the whole of that space. If we then position two of our snooker balls at those corresponding points within our symmetrical universe, then it remains a theoretical possibility that two distinct particulars nevertheless are identical in all their non-relational and relational properties. (This sounds a bit complicated, but Figure 1 shows what's meant.) On the bundle theory, they again collapse into each other.

This is a complicated argument. A short summary might help. We tried to separate indistinguishable particulars on the basis of them having different locations. But either those locations are themselves particulars, in which case we have not succeeded in eliminating particulars, or locations are just distinguished by their relations to each other. And in the latter case, the possibility of a symmetrical structure means that we could have
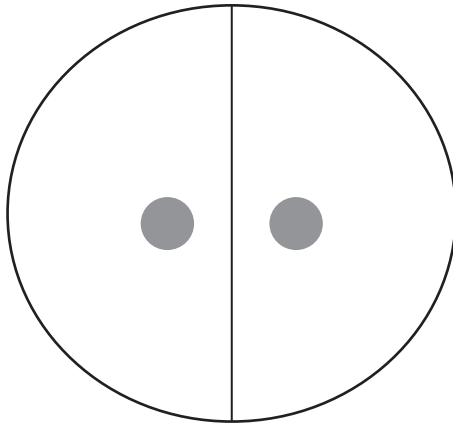
1. A symmetrical universe

two particulars that were not distinguishable even on the basis of location.

What we just had was a proposed first way for the bundle theorist to avoid the implication that particulars with all the same properties collapse into one. As that didn't seem to work, here is a second proposal. The objection, that the theory entails bundles with all the same properties must be one and the same, strikes only if the properties are to be understood in a certain way: as nothing like particulars. But there are other conceptions, as we will see in Chapter 2. Perhaps those properties are particularized in some way. Hence, the red in this bundle might be a different thing or instance from the red in another bundle. There might then be the possibility that there are distinct particulars with all the same properties. They consist of the same types of property but different instances of them. Isn't this what we think of all the red snooker balls? The red of this ball is not the same as the red of that ball. They are two different instances of red.

But there is again a problem with this apparent solution. We have saved the bundle theory but at a cost. An advantage of the bundle theory, it was noted, was that it accounted for particulars entirely in terms of properties. Particularity was reduced away in terms of properties. But it now seems that we are able to salvage the bundle theory from the objection that two identical bundles would collapse into one only if we understand properties in some way as particulars. We spoke of having two distinct instances of red and a property instance looks like some kind of particular. So to make our bundles behave more like the particulars that we take objects to be, we have had to make our properties like particulars. Particularity has managed to sneak back into the theory.

There are countless mistakes that we may have made along the way. But it looks like we might be forced to conclude that particularity is an irreducible feature of reality, for there could, in

Metaphysics

theory, be two distinct particulars whose distinctness did not consist in them having different properties.

So what, then, is a table? After the considerations in this chapter, it seems that we have to say it is a particular that bears certain properties but is not identical with, nor reducible to, those properties. The table was chosen arbitrarily as the object we examined, and it thus seems safe to generalize from it. We should then give the same answer for any other object.

The properties of particulars have been mentioned throughout this chapter. We need next to consider what these things are supposed to be, if indeed they are things at all. We move on, therefore, to this topic.

What is a table?