Introduction

In general, contemporary metaphysics is deeply sceptical of the familiar objects in which common sense believes. It is far more ready to attribute reality to entities that are much smaller – to the particles and wave packets and strings which microphysics treats as real, or to the "mereological simples" for which philosophical reflection provides some support. Any such view must find some way of explaining why there appear to be familiar medium-sized objects in the world. Many metaphysicians suppose that we can do just that. We can explain why it appears that the microparticles of the world compose familiar objects, why it appears that these objects persist across careers in which they lose and gain component microparticles, and why it appears that these objects have and exercise causal powers. The main business of this book is to argue that leading examples of such reductive explanations fail. For time and again such explanations project downwards, onto the small entities of the preferred ontology, structures and relations and features that properly belong to familiar objects. Such projection is harmless so long as one allows that there also are, in addition to the small entities, the familiar objects that form the starting point of the projection. But if – as is generally the case – the aim is to expunge familiar objects from ontology, the invocation of such structures and relations and features is illegitimate. The opponents of familiar objects are then helping themselves to shadows cast downwards, onto the level of the preferred small entities, while denying that the sources of these shadows exist.

The metaphysical position which this book is intended to support is that at least many of the familiar objects that common sense recognizes are mind-independently real. The book begins with a chapter that undertakes to refute two false friends of this commonsense ontology. The first is the modal conventionalist, who holds

that the general ways in which nature's kinds are marked out from one another, and the general ways in which persistences of members of those kinds are delimited, are fixed by the "descriptive content" and "referential intentions" that we associate with our sortals and matter-names. Such a view fails to treat as mind-independently real the phenomena of sameness in kind and of persistence across time. It therefore falls short of *realism* about familiar objects (indeed about any objects) since it makes the natures that objects share with others of their kind, and the careers which each individually traces out, be functions of our cognitive and linguistic practices. The other false friend is the explosivist, who is happy to award mind-independent reality to the familiar objects of common sense, but who cheapens that award – indeed nullifies it – by awarding reality likewise to every imaginable crosscutting of the world's individuals and kinds. The modal conventionalist thinks that nothing is required of the world in order for our general ways of tracing persistences and detecting kind-samenesses to be correct; the explosivist thinks that nothing *special* is required, since any general ways of doing this cannot fail to track real persistings and real samenesses in kind. Chapter 1 argues that neither conventionalism nor explosivism embodies an adequate understanding of the ways that our talk about sameness in kind, and about numerical persistence across time, functions in our cognitive economy.

Chapter 2 deals further with modal conventionalism. Of the two false friends of common-sense ontology, conventionalism and explosivism, conventionalism has been the more influential, and so deserves the more protracted treatment. One form which conventionalism has assumed is a view that might aptly be labeled "ontological relativism." Chapter 2 examines one arresting argument for ontological relativism, and contends that that argument fails. More radically, chapter 2 argues that any argument for ontological relativism must fail – that the view is conceptually untenable.

Between them, chapters 1 and 2 raise serious objections against the main currently prevailing forms of antirealism about material objects. Chapter 3 then sets forth a realist position on material objects. It articulates the connections between realism about the *existence* of material objects and realism about the two forms of *sameness* discussed in chapter 1, and shows what a realist ontology that incorporates all three elements must look like.

The book then turns to the opponents of common-sense ontology. As puzzles such as the ship of Theseus have made vivid for millennia, familiar objects, if real at all, seem to survive across a messy variety of alterations – and not just of loss and gain of component particles, but of hard-to-delimit qualitative alteration. Here the desire for a cleaner ontology may motivate an ontological preference for entities that are *temporally* smaller than familiar objects appear to be – for temporal stages, each of which lasts no longer than the shortest possible physical change. Such temporally tiny entities can explain away the appearance that there are familiar objects that persist over long careers by serving as the truth-makers for claims expressing that appearance. If a temporal stage of the right qualitative character stands in temporal counterpart relations to other stages having the right qualitative character, then, say stage theorists, a sentence that asserts or presupposes the persisting of a familiar object can be rendered true. But I argue that temporal counterpart relations – if they are not going to saddle us with an explosivist account of the world's persistences – constitute an illegitimate projection downward from the careers and powers of familiar objects.

The second group of opponents are the causal exclusionists. From the time of Plato's Eleatic stranger it has seemed plausible that familiar objects, if real, must be capable of bringing about effects. But any familiar object is wholly composed of entities that, spatially, are vastly smaller – in particular, the microparticles of physics. If physics is closed and complete, it may seem, then, that the several doings of the component microparticles must between them cause anything which the familiar object may be said to cause. We have apparent overdetermination, which is apparently intolerable, and the apparent victory goes to the microparticles. But to which microparticles? When, I shall argue, microparticles are grouped in the ways relevant for awarding them efficacy over the effects that common-sense attributes to familiar objects, they are grouped in ways that illegitimately project downwards, from the level of the familiar objects themselves. Causal exclusion arguments, focused on familiar objects, generally fail. They fail in particular for the apparent exclusion of mental causation. Mental causation is the case I shall use to focus the debate about whether a familiar object's microparticles steal away the apparent efficacy of that familiar object. That is, I shall discuss whether beliefs and desires – states of that most familiar of familiar

objects, a person – genuinely cause behavioral outcomes. But I will indicate how the argument generalizes to the cases of other familiar objects.

The third group of opponents are the sceptics about composition. It seems secure, to these opponents, that tiny entities entirely occupy any volume in which common sense supposes a familiar object to be present. But is there also the one large object which these entities seem to compose? These philosophers begin with uncertainty about what, in general, composition might amount to, and proceed to scepticism about whether there objectively is any such phenomenon at all. It unquestionably appears that microparticles compose such (relatively) large objects as dogs and trees and desks – but that appearance may amount to no more than that the microparticles themselves are "dogwise" (or "treewise" or "deskwise") arranged. I argue that such adverbial arrangements are a projection downwards from familiar objects. They are, indeed, perfectly real *if* the objects from which the projection proceeds are themselves real. That is, there is a genuine phenomenon of microparticles' being dogwise or treewise or deskwise arranged, if this phenomenon just amounts to the fact that those microparticles jointly occupy (and are confined to) the entire volume in which a dog or tree or desk exists. If that sort of fact obtains, indeed, it likewise provides the analysis of what it is for those microparticles to *compose* a dog or tree or desk. But the opponents in this third group want to explain away, rather than affirm, the reality of such familiar objects. And if familiar objects do not really exist, the phenomena of dogwise or deskwise arrangedness are purely imaginary.

A fourth group of opponents thinks that composition is not something which microparticles owe to the reality of the familiar objects within which they are found, but something that they possess in their own right. Indeed any plurality of entities whatever composes something, says this fourth group. Composition is a "free lunch" (in Armstrong's phrase), which comes automatically with the bare existence of the components. This is the doctrine of universal mereological composition (UMC). The proponent of this doctrine qualifies as a third false friend of the ontology of common sense. The proponent appears to be a *friend* because, at any moment at which common sense supposes a familiar object to exist, she will find a mereological sum of microparticles that occupies just that volume which common sense supposes the familiar object to occupy (with the small qualification

that the mereological sum will have relatively precise boundaries, while the familiar object apparently has vague boundaries). But the proponent of UMC is a *false* friend, not just because of her explosivist commitments, but because, as I shall argue, her stand-ins for familiar objects are compositionally brittle, while familiar objects themselves are compositionally flexible. That is, across the phases of its existence, a familiar object might have incorporated different microparticles from those that it did; not so, I argue, the mereological sum of microparticles that is located where that familiar object is. UMC is a false doctrine, I shall argue. The composed objects which *it* countenances would in general be characterized only by certain structural properties, properties that fail to contrast to greater and lesser degree with their own proper contraries. But determinate contrast-with-contraries is constitutive of the very identity of any genuine property.

The argument of this book is defensive. The book identifies inadequacies in contemporary attempts to "explain away" familiar objects, attempts intended to show that *no* familiar objects really exist. Some readers may find themselves wishing to see positive arguments *in favor of* familiar objects. As responses to the contemporary opponents of familiar objects, these would be arguments to the effect that *some* familiar objects are perfectly real – not necessarily that every familiar object posited by common sense, or by one of the special sciences, is real. I will offer no argument for that more limited conclusion, because I believe that any such argument would be question-begging: it would have to proceed from premises that assume that at least some familiar objects exist. For the proponent of familiar objects, as I see matters, the situation is exactly that of Neurath's boat. We can suspect individual planks of rot, can remove them and examine them, and can even replace them if need be. But we do this while afloat on the boat – while standing on other planks. There can be no systematic justification for standing-on-planks-in-general-and-as-such.

The history of post-Cartesian philosophy of course contains many efforts at establishing the reality of familiar objects – objects such as trees and dogs, stars and cells, perhaps even desks and pencils. These arguments all proceed from premises, allegedly more secure than the conclusions to be established: that *we* exist, and engage in various cognitive and perhaps practical activities. But if we ourselves are familiar objects – human organisms, say – and if our featured

activities engage and are directed at other familiar objects, then these arguments are question-begging in just the way I have indicated. The only alternative is to start from a picture of ourselves as transcendental egos, and to secure the existence of familiar objects by virtue of their relation to transcendental mental activity. This runs counter to the naturalist position, to which I subscribe, that we are objects *in* the world of familiar objects, distinguished mainly by the history of natural selection that has fashioned us. Beyond that, the only reality which such an argument can deduce for familiar objects is a mind-conferred, mind-dependent reality.

Not here. The thesis of this book is that familiar objects – at least some of them – are mind-independently real. Their detractors seek to impugn the ontological status of familiar objects by using shadows which those objects cast, while denying that the shadows have a source.

1

Two false friends of an ontology of familiar objects

Judgements asserting one or another of two kinds of sameness are crucial, I shall argue, both for our practical mastery of the world and for our theoretical understanding of it. On the one hand, there are judgements saying that one object is the same in kind as other objects, or that some matter is the same in kind as matter found elsewhere. On the other hand, there are judgements saying that the object in front of us is numerically the same object, or that the matter is the very same matter, as we encountered earlier or will encounter later.

In making these judgements we call upon observation and understanding. In order to affirm sameness in kind, we must observe that various similarities obtain between one object and others, or between matter here and matter elsewhere. In order to affirm persistence across a single episode of observation, we must observe that an object (or some matter) has moved continuously, while retaining largely the same features, and in order to affirm persistence across separate episodes, we must observe that the object (or the matter) now before us presents features appropriately related to those observed in an object (or some matter) encountered at other times. But we must also understand *which sorts* of similarities indicate sameness in kind, *which sorts* of relations mark out persistences. We must understand, for example, that sameness in chemical microstructure indicates sameness in kind, as between two portions of matter, while sameness in color does not, and neither do sameness in heft or in location. We must understand that specific sorts of sameness or change are to be expected in a persisting object of the kind to which the object observed earlier and the object observed later belong, and specific spatiotemporal relations to the place of the earlier observation.

To speak of a cognitive performance as a case of understanding is to say that it can fail as well as succeed: where understandings occur, misunderstandings are possible. Just what is required of the world, in

order for the understandings that undergird our judgements of kind-sameness and numerical persistence to succeed? What are the truth conditions for our implicit ideas about which *sorts* of properties are constitutive of kind-sameness and which *sorts* of relations indicative of numerical persistence?

To these questions, prevailing philosophical opinion offers one or another of two markedly puzzling answers. One is that *nothing* is required of the world in order for our ideas about how nature's kinds are delimited, and how nature's persistences are marked out, to be true – our ideas are true, but they do not have truth conditions. For the sentences that express these ideas are analytic; they are not true in virtue of the world's being one way or another.[1] This is the claim of the "modal conventionalists" (or "modal conceptualists"), to whose views I return in section 1.5. The other prevailing opinion is that nothing *special* is required for these ideas of ours to be true. They are true in virtue of ways the world is, but the world is many ways at once. It contains crisscrossing kinds so numerous that any scheme for assigning objects to kinds is bound to capture a way the world is.[2] It contains vastly many colocated objects which persist over quite different spans of time – objects which may even persist in quite different forms from one another, or at quite different locations – so that any scheme we might subscribe to, for tracing persistences, is bound to be right.[3] These are claims made by philosophers who might all appropriately be called "explosivists," to whose views I return in section 1.4.

Either sort of answer is markedly puzzling. If (as I shall argue) our judgements of kind-sameness and numerical persistence are crucial both to our practice and to our theory, it is nothing short of amazing that whatever general scheme we might embrace for making such judgements is bound to succeed – that the only error possible is error in executing that scheme, by making faulty observations. The business of this chapter is to articulate a realist picture of our implicit

[1] Thomasson, *Ordinary Objects*, pp. 67–68; Sidelle, *Necessity, Essence, and Individuation*, p. 128.

[2] For an illustration, see Hirsch, *Dividing Reality*, pp. 24–25 – but concerning Hirsch's own position, see note 21 of this chapter.

[3] Hawthorne, *Metaphysical Essays*, p. vii; Sosa, "Existential Relativity," p. 142; Sider, *Four-Dimensionalism*, p. 133.

ideas about how kind-samenesses are constituted and how persistences are marked out. On this realist picture, these ideas do indeed have truth conditions. The modal conventionalists are wrong. But not every such idea that we might embrace will be true – the truth conditions are sparse. Explosivists are wrong. I shall *articulate* a realist picture of our schemes for tracing kind-sameness and numerical persistence, but will not offer much in the way of positive argument for that picture. For, in my opinion, the presumption must be that *some* realist account of these schemes is correct; it is simply not credible that on matters so crucial for our survival and our understanding of the world, just *any* way that we might proceed cannot end up steering us wrong. But I shall take pains to explain how, *if* the realist picture I offer is correct, both conventionalism and explosivism misunderstand our judgements about kind-sameness and numerical persistence.

1.1 THE POINT OF AFFIRMING KIND-SAMENESS

Why does it matter to us to note that two (or more) objects are the same in kind, or that some matter here is the same in kind as the matter over there? The standard answer is that such judgements guide inductive inferences. That is the answer I shall set forth in this section, adding some not entirely standard details.

The kinds of the world typically count for us as falling into families: there are kinds of animals, kinds of food, kinds of fluids and plants and even artifacts. Typically, the kinds within any such family count for us as collectively characterized by certain *sorts* of properties, and as individually characterized by just one instance of these sorts. Thus we suppose that each kind of animal is characterized by a particular style of locomotion, a particular diet, and a particular (if roughly defined) body shape and size. Foods of the various kinds are, in general, each characterized by an aroma, a texture, a taste, and some even by a particular way we feel in response to eating them. Different kinds of fluids are, in general, each characterized by a particular color and scent and viscosity, and perhaps even by a particular boiling or freezing point. Each kind of plant has a characteristic pattern of growth, and a characteristic morphology when mature. The various kinds of artifacts typically count for us as each having a particular shape and design, a particular use, and as having

a characteristic association with a setting *in* which or *with* which that use gets performed.[4]

To put it generally: the kinds in each family count for us as selecting just one out of the properties in each of several (or even many) ranges of properties. Each kind does this, we suppose, in member after member or portion after portion. For this reason, judgements that a plurality of objects (or of portions of matter) belong alike to a common kind have the function of directing certain kinds of inductions. The judgements say, so to speak: you need only discover which member out of each relevant property-range is present in a few observed members of a kind, in order to be sure that that same property is present in other members of that kind (or other portions, wherever located). Judgements of kind-sameness enable *instantaneous discoveries about items that are remote*. Once we have made such discoveries, moreover, judgements that newly encountered objects belong to that same natural kind enable us to say: these newly encountered objects have properties that may not yet have been revealed to our observation. We need only observe *enough* about the newly encountered objects to judge that they belong to that same kind, in order to be sure that these newly encountered objects have other properties beyond those we immediately observe – including properties that we *cannot* detect in the current observational setting. Judgements of kind-sameness then enable *amplified observation of items that are present*.

How do we know just which *sorts* of properties it is, from which each kind in a given family selects just one instance? Modal conventionalists maintain that this knowledge is a priori. It is analytic, they say, that the animals in a particular natural kind are all characterized by a particular body shape and a roughly defined body weight upon maturity,[5] analytic that food of a given kind is everywhere characterized by a particular taste and texture and smell. I shall return to this position in section 1.5. For now I merely note that it seems implausible that knowledge of the sort in question is *never* empirical. It was an empirical achievement of some moment, one might well suppose, to learn that the various fluids and

[4] Elder, "On the Place of Artifacts in Ontology."

[5] Thomasson, *Ordinary Objects*, pp. 38–44 and 48–53; Sidelle, *Necessity, Essence, and Individuation*, p. 112.