1

CLASSIFICATION

1.1 Introduction

1.1 This chapter focuses on the common-place activity of distinguishing and classifying objects of various kinds. You might wonder: Why start a book on metaphysics with a discussion of classification?

1.2 There are a couple of reasons. First, lurking behind this common-place activity are a lot of metaphysical puzzles and questions! One of the cool things you'll discover as you study metaphysics is that the world is a lot more complicated and much stranger than you initially might have thought. A good way to illustrate this is to start with something down-to-earth and rooted in our ordinary ways of thinking and talking. Once you see that even something that is seemingly straightforward has a tangle of puzzles hiding behind it, you'll start to suspect that philosophical perplexities can arise about pretty much anything.

1.3 Second, the metaphysics of classification will provide a nice springboard for the discussions to follow on the metaphysics of properties (in Chapter 2) and the metaphysics of parts and wholes (in Chapter 3). These parts of metaphysics are somewhat more abstract than the more down-to-earth things we'll begin with here, but they are intimately related to the metaphysics of classification, as we will see later on.

1.4 Let me give you a breakdown of this chapter. In Section 1.2, I will introduce and explain a distinction between two different ways of classifying objects: an objective and a subjective way. In Section 1.3,

This Is Metaphysics: An Introduction, First Edition. Kris McDaniel.

I will discuss some cases in which it seems that we have mistakenly taken a merely subjective classification to be an objective one. But even if we sometimes do make this sort of mistake, it does seem like we still often succeed in objectively classifying objects. Section 1.4 will present an argument for the conclusion that some things do objectively belong to each other. In Section 1.5, we will explore the question of what it takes for things to objectively belong together. This will naturally lead us to a discussion of the connection between the metaphysics of classification and the metaphysics of properties in Section 1.6. (And Chapter 2 will be focused more generally on the metaphysics of properties.) Finally, in Section 1.7, we'll close with some further questions about classification to consider.

1.2 Two Kinds of Classification

Let's start with something that seems easy. Think about this list of things: a cat, a dog, a kangaroo, a fish, and a loaf of bread. Suppose you were asked, "Which one of these things *does not belong* with the others?" I don't think you'd have a problem answering. You'd unhesitatingly single out the loaf of bread. This wouldn't even be a hard question for a small child. My five-year old unhesitatingly singled out the loaf of bread too. 1.5

Suppose we take out the loaf of bread from the list and ask again, "Which one of these things *does not belong* with the others?" You might struggle a bit more this time, but you'd probably exclude the fish, although that isn't the only defensible answer. For example, you might exclude the fish because each of the remaining three is a mammal. However, you might instead exclude the kangaroo on the grounds that each of the remaining three is commonly taken as a pet. (Even in Australia, it is not very common for someone to have a pet kangaroo.) 1.6

But suppose I gave you the following list of things: a neutron star, the number 2, the dream you had last night, and a blade of grass. Now the question "Which one of these *does not belong* with the others?" is much harder to answer. Why is this? I suggest that it is because any way of excluding one of these items from the list leaves us with a list of three things that don't really belong together any more than the original four. And you recognize, at least implicitly, this fact. 1.7

But why do some groups of things belong together while other groups of things do not? A complete answer to this question requires some metaphysics: 1.8

specifically, we need a theory that explains what belonging together amounts to in general. This theory will be a metaphysical theory of classification. (We'll discuss this further in the Section 1.5.)

1.9 A classification of a group of things is just a way of breaking up that group of things into groups. Let's say that a classification of things is a *good* way of classifying things when it breaks things into groups, and anything in one of those groups belongs with all the other things in that group but doesn't belong with the things that are in a different group. Venn diagrams provide a good way of illustrating this idea.[1] Suppose we have a group of things *w*, *x*, *y*, and *z*.

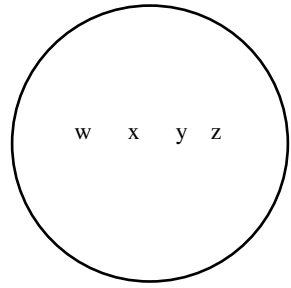

Suppose that *w* and *x* belong together and that *y* and *z* belong together, but also that no collection of exactly three of them belong together. Then a good classification would divide up our initial group of four things into two groups of two things. We could represent this classification with a picture:

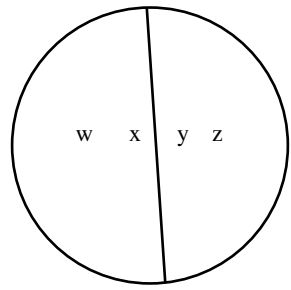

In this classification *w* and *x* have been put together and separated from *y* and *z*, and *y* and *z* have also been put together.

[1] http://www.mathsisfun.com/sets/venn-diagrams.html

So far, we have kept things pretty simple. One complication is that belonging together probably comes in grades or degrees. That is, some things will belong together to a greater extent than other things. Think about this list of things: me, my wife, my two daughters, and my dog Ranger. Ranger is the odd man (odd dog?) out in this list. But now consider a larger list: me, my wife, my daughters, my dog Ranger, and the number 2. Despite being even, the number 2 is the odd number out. What this shows is that although Ranger doesn't belong in the first group as much as my daughters do, he definitely belongs in the second group much more than the number 2. So, he must belong with me, my wife, and my daughters to some extent. 1.10

What this means is that instead of focusing simply on the question of whether a specific classification is a good way of classifying, we should also consider the more general question of when a given classification is a *better* classification than another. We won't have a complete theory of classification if we don't consider this more general question. 1.11

The second complication is that there are different reasons why we judge that things belong together, and these different reasons will sometimes lead to different and apparently conflicting ways of classifying objects. But these different ways of classifying objects don't really conflict. We already saw an illustration of this phenomenon. Let's return to one of the lists of things that we thought about at the beginning, specifically the one that consisted of a cat, a dog, a kangaroo, and a fish. There seemed to be two equally respectable ways of breaking this list into groups. We might separate the fish from the remaining three on the grounds that the remaining three are mammals; but we may also separate the kangaroo from the remaining three on the grounds that only the remaining ones are commonly taken as pets. Is there any point in trying to decide which of these two classifications is *better*? Isn't how we classify things largely dependent on what we are interested in, what we care about, what we desire, and so forth? 1.12

Sort of. There's an important difference between the two kinds of classifications just mentioned. One of them classifies animals by looking to see whether they are related to *us* in some interesting *social* way: the classification that excludes kangaroos is based on the observations that *we* don't interact with kangaroos in (some of) the ways in which we interact with dogs, cats, and fish. But it wasn't inevitable that we don't typically have kangaroos as pets. Imagine a possible situation in which we have kangaroos as pets instead of cats: if that possible situation had actually happened, then we probably would have felt that it was cats rather than kangaroos that do not belong on the list. 1.13

1.14 But now think about the other kind of classification, the one that excludes fish because they are not mammals. This classification isn't based on our desires or interests, and it doesn't classify things on the basis of how these creatures relate to *us*. Mammals could have existed even if no human beings had ever evolved on the planet. And mammals still would have belonged with each other in a way that cats, dogs, fish, and kangaroos do not. Moreover, we can't easily imagine possible scenarios in which, for example, dogs fail to be mammals. We can imagine possible situations in which there are things that have the same outward physical appearance as dogs and that have similar behaviors to dogs, which aren't themselves mammals. But that's not the same thing as imagining a situation in which dogs aren't mammals.

1.15 One lesson to draw from these observations is that, roughly, there can be (at least) two kinds of systems of classification. We might classify objects because we are interested in how they relate to *us*. We can think of these as *subjective classification systems* because they are driven by our interests, desires, values, and so on. But a system of classification might also classify things on a more objective basis, that is, on grounds that are independent of whether that system of classification reflects anything about how reality relates to us or what we care about. Some things *objectively belong together*. An objective classification system is one that classifies things together in a way that matches how they objectively belong together.

1.3 Classification Confusions

1.16 Sometimes we make mistakes about what things objectively belong together, and these mistakes are discovered only after painstaking scientific investigation. Here are some relatively straightforward examples to consider. An ordinary culinary classification might lump carrots, beets, and tomatoes into one group, while lumping oranges, apples, and grapes into another group. But from a more scientific perspective, tomatoes objectively belong more with oranges, apples, and grapes than they do carrots and beets. And as awareness of this fact has spread, how people use language to classify things has changed.[2] Two hundred years ago the sentence "Tomatoes are fruits" might have seemed

[2] Whether tomatoes are fruits might seem completely trivial, but interestingly this is a question that has been addressed by no less than the Supreme Court of the UnitedStates!http://www.npr.org/blogs/money/2013/12/26/256586055/when-the-supreme-court-decided-tomatoes-were-vegetables

like a ridiculous thing to say to your average speaker of English, but that's not the case today. Now many of us are happy to utter, "Tomatoes are fruits," and believe that we say something true when we do. There is an interesting question of whether the word "fruit" has changed its meaning over time as a result of a greater awareness of the fact that tomatoes objectively belong more with grapes than they do with carrots. If the word has changed its meaning, then what was said two hundred years ago by "Tomatoes are not fruit" might well have been true too! But it would still have been false that tomatoes objectively belong with carrots more than they do with grapes.

A similar observation can be made about the word "fish." Are whales 1.17
fish? They used to be classified with other things that were called "fish," but now we know that whales do not objectively belong with fish to the extent they were once thought to, and as a consequence, there is pressure not to call whales "fish."[3] One of the things that such scientific discoveries can teach us is that our beliefs about which things objectively belong together are subject to serious revision. An even more recent case is the reclassification of Pluto as not being a planet.[4]

Social scientists—psychologists, sociologists, anthropologists, and so 1.18
on.—would be good people to consult if we wanted to know why people lump things together for subjective reasons. Maybe biology, zoology, animal psychology, and other sciences like them, might also play a role in helping us figure this out. And it will sometimes be a tough job. Occasionally the explanation for why a group of people ended up classifying some objects as belonging together can be so subtle that the people doing the classifying aren't even aware of these reasons. And when this happens, sometimes people falsely believe that the things that they have subjectively lumped together really objectively belong together. As we will see in a moment, at times this can also affect how people classify each other.

Consider, for example, the practice of classifying persons on the basis of 1.19
race. Many scientists believe that racial classifications are not backed up by the biological facts.[5] Consider, for example, how a parent classified as belonging to one race can have a child classified as belonging to another

[3] If you are curious about why whales are not fish: http://scienceline.ucsb.edu/getkey.php?key=2536

[4] Here's what NASA has to say about Pluto: http://www.nasa.gov/audience/forstudents/k-4/stories/what-is-pluto-k4.html#.U6nCQ_ldWSo

[5] The concept of race is in need of retirement, according to Professor Jablonski. http://edge.org/responses/what-scientific-idea-is-ready-for-retirement

race. The parent and child will likely have more biologically in common with each other than they would with other members of "the same race." But as any student of the history of racism is aware, people discriminated on the basis of race because they thought that "racial classifications" grouped people into collections that objectively belonged together.[6]

1.20 If this is correct, then racial classifications are very subjective in the sense described earlier: racial classifications are based on contingent features of societies rather than on genuinely important objective features of the world. This seems to be a view that most people who accept the idea that race is socially constructed accept, though many augment this basic position with other philosophical theses.

1.21 Some philosophers have argued for an even stronger conclusion: there are no races. This is the position defended by the philosophers Kwame Anthony Appiah and Naomi Zack, among others.[7] We'll call this position *race eliminativism*. Here are the basic ideas behind race eliminativism: not all attempts at classifying entities are guaranteed to succeed. Sometimes we introduce terms to stand for groups that we initially take to objectively belong together—but because they don't objectively belong together, we actually fail to classify things with those terms. Here's an uncontroversial example of a failed attempt to classify people. A while ago, certain women were thought to have magical powers and were persecuted as a consequence. These persecuted women were called "witches." According to their persecutors, not all women are witches—just those who have important characteristics in common, specifically magical abilities. The people who labeled these women "witches" thought that the women they persecuted objectively belonged together. But since there aren't magical powers, the persecutors' attempt to

---

[6] The *Stanford Encyclopedia of Philosophy* is a wonderful resource. It's filled with excellent introductions. Here is the one for "race," which I highly recommend: http://plato.stanford.edu/entries/race/

[7] Appiah is currently a professor of philosophy and law at New York University. His website is: http://www.appiah.net Most of his published papers are available there, including: http://appiah.net/wp-content/uploads/2010/10/The-Uncompleted-Argument-Du-Bois-and-the-Illusion-of-Race.-Critical-Inquiry-12.1.-1985.pdf This paper discusses Appiah's views on the metaphysics of race.

Zack is currently a professor of philosophy at the University of Oregon. Zack's views on the metaphysics of race can be found in her book *Philosophy of Science and Race* (London: Routledge, 2002). Her faculty profile is here: http://philosophy.uoregon.edu/profile/nzack/

classify these women together was a colossal failure. The lesson we learn from their failure is that while there might be women who were thought to be witches, there really are no witches. Similarly, according to race eliminativism, when words that allegedly refer to races were first introduced, the people who introduced these words thought that there were really important and deep biological differences between human beings, and on the alleged basis of these differences, they attempted to objectively classify people. But these differences have proven to be either non-existent or merely superficial. So, the attempt to objectively classify individuals on the basis of race is a failure just like the attempt to objectively classify individuals as witches. And so, according to race eliminativism, just as we learned that there are no witches, we should also conclude that there are no races.

1.22 One of the key theses of race eliminativism is that classification by race is *intended* to be an objective classification. And since this intent failed, the attempt to classify via race fails too. And so, there are no races. But if race eliminativism is true and there are no races, why do people talk as if there are? It is important to keep in mind the distinction I noted earlier, between two sorts of classification: one that is based on whether things objectively belong together and the other that is based on more subjective reasons. Probably the full explanation for the prevalence of classifying people on the basis of "race" will be understood only once we have better grappled with both the historical and contemporary reality of various forms of racial stratification and oppression. The fact that we do classify for subjective reasons might generate the appearance that there are races even though, according to race eliminativism, there really aren't any.

1.23 This is not to say that there are no biological bases at all for apparent racial differences. There is a biological explanation for why, for example, two people have different skin colors. But the existence of these sorts of explanations doesn't mean that classification on the basis of races is particularly objective. Here's an analogy to consider. Suppose that instead of thinking of people in terms of "races" we thought of people in terms of "heights." People under 5' tall were called "the lows"; people 5' to 5'2" were called the "low mids"; people 5'2" to 5'6" were called "the mids"; and so on, and so on. Presume, in this other possible culture, that what height you belonged to mattered as much as race sometimes seems to matter in our culture. (Suppose, for example, that there was systematic discrimination against the lows.) Clearly, these classifications of people in terms of height are much more subjective than objective, even though in each case there is going to be a biological explanation for why a person is of a given height.

1.24 Race eliminativism isn't the only position in the philosophy of race. It's a rich and varied field, like all areas of philosophy! Recently, the view that races are biologically real has been defended in a series of interesting papers by Quayshawn Spencer, who argues that, at least in the United States, racial classifications correspond to five global population groups corresponding to genetic clusters.[8] On Spencer's view, dividing the human population into these five groups results in genetic clusters that minimize genetic differences among individuals within a group but maximize genetic differences among individuals from different groups. Spencer clearly notes, however, that these genetic differences do not imply any moral or intellectual differences between the groups. (There are some suggested readings at the end of this chapter for those of who you wish to pursue topics in the philosophy of race more.)

1.25 Let me briefly sum up what has been discussed so far. We've explored two kinds of classification systems: those that classify on the basis of our subjective interests, feelings, and whatnot, and those that attempt to classify things on the basis of whether the things in question objectively go together. We've seen that we can make mistakes about whether some things objectively go together, and that some of these mistakes can even be very harmful. We discussed classification by race as an illustration of one way this might happen.

1.26 But the fact that we can and often do make mistakes about what things objectively go together shouldn't by itself automatically make us skeptical that nothing really objectively belongs with anything else. In order to get us to be skeptical across the board about the possibility of objectively classifying entities, we'd need a much more powerful argument than one that begins with the premise that we often screw it up (and that we are really bad at it when it comes to classifying people). On the other hand, it would be nice to have a positive argument for the conclusion that some things really do objectively belong together. We will discuss one such argument in the following section.

1.4 Do Things Objectively Belong Together?

1.27 One of the most obvious ways in which we classify things is by using language. We make use of general terms, such as "dog," "cat," and "table" to group things together under common headings. Once we've done this, we

[8] Quayshawn Spencer is currently a philosopher at the University of Pennsylvania. His webpage, which contains many of his published papers, can be found here: https://sites.google.com/site/qnjspencer/publications

can use these terms in sentences that communicate very general facts, such as that dogs typically are furry and friendly. In addition to terms like "dog" and "cat" that seem to stand for kinds of things, we can also classify things by way of predicates like "red" and "circular."

1.28 But there are some interesting apparent differences: dogs do seem to form a kind of thing in a way that, say, red things do not. Suppose I tell you that I have a red thing in a box: it would be reasonable to ask what kind of thing it is, and if you did ask this question, and I replied by saying that it is a red thing, or a colored thing, you'd probably think that my answer was pretty weird. Suppose I tell you that I have a dog in a box: it would be a little odd to ask what kind of *thing* is in my box, though it wouldn't be at all odd to ask what kind of *dog* is in my box. (You should probably also ask *why* I put a poor dog in a box, but since this is a book about metaphysics rather than ethics, we won't pursue that question further.) We use words like "dog," "cat," "car," "desk," "proton," and "human being" to stand for kinds of things, and we use words like "red," "tall," "solid," "fast," and "furry" to stand for features of things. There might be an important metaphysical difference between kinds of things and features of things. But there is also an important similarity that is relevant to our discussion of classification: we use both of these types of expressions to mark differences in things: these things are the things that are red, rather than those things that are some other color, such as green; these are things that are dogs rather than those things that are some other kind of animal.

1.29 Reflecting on the various ways we use language to classify objects can open us up to new philosophical problems. Here is an interesting thought experiment suggested by Eli Hirsch.[9] Suppose we were to encounter a community of human beings who speak a language that is in many respects like English but with some significant differences. First of all, these people don't have words for colors like "red," "purple," "green," and so forth in their language, and they don't even have any words that are synonymous with those words. They also don't have words for shapes, such as "circular," "triangular," and so forth. Although these kinds of words are missing from their language, this is not because the people in this community differ from us in

[9] Eli Hirsch is currently a professor of philosophy at Brandeis University: http://www.brandeis.edu/facultyguide/person.html?emplid=3b8c58f368269c9b2a81bea2caf5cd32e3d4e84a An important book by Hirsch on the subject of classification (among other metaphysical topics) is *Dividing Reality* (New York: Oxford University Press, 1993).

what they can perceive with their eyes. (I suppose that if human beings had never evolved with the ability to see things, we wouldn't have had words for colors either.) Their language does contain some words that English doesn't contain. For example, it contains the word "gricular," which stands for a feature that things have when they are *either* green *or* circular, and it contains the word "grincular," which stands for a feature that things have when they are *either* green *or not* circular, and it contains the word "ngricular," which stands for a feature that things have when they are *either not* green *or* are circular. Let's call this language "the gruesome language."

1.30 The gruesome language is peculiar. I'll say a few more things about the gruesome language that will probably amplify that impression. There are some words in English that can be explicitly defined. A stock example is the word "bachelor," which can be explicitly defined as "someone who is an unmarried, eligible to be married, adult human male." Competent users of words like "bachelor," which is a word that is capable of being explicitly defined, should be able to at least gesture at its definition. If a person whose first language is not English asks you what "bachelor" means, it would be appropriate to provide her with the abovementioned definition. But not every word in English is capable of being explicitly defined, even by fully competent speakers of English. When you give an explicit definition of a word, obviously and unavoidably you do this by using other words. Some of these words might themselves be capable of being explicitly defined, but at some point the chain of definitions has to come to an end. There must be some words in English that aren't capable of being explicitly defined, even by competent users of English. I suspect that words like "red" and "green" are good examples of words that we can't explicitly define. Suppose a person whose first language is not English were to ask you what "green" meant. You could do the following things to help her out. You could show her samples of things that are green, such as American dollar bills, or the color of grass or leaves in late spring, or if you are especially fortunate, some lovely emeralds. In other words, you would help this person learn what "green" meant by *pointing at things that happen to be green*. That's a very different activity from giving an explicit definition of the word "green." I suppose that you could also say that green is the color that things have when they reflect light with a wavelength of around 510 nanometers, but you shouldn't think that by doing this you are giving a *definition* of "green." Definitions of words are the sorts of things that lexicographers discover, but that green things reflect light of a wavelength in the neighborhood of 510 nanometers is a discovery of physics. Perfectly competent speakers of English hundreds of years ago knew what "green" meant

but most of them knew next to nothing about the physics of light. (That's true of a lot of competent speakers of English today too.) Competent speakers of English can't give an explicit definition of "green."

The speakers of the gruesome language are in a similar position with respect to the words "gricular," "ngricular," and "grincular." If you asked them what they meant by "gricular," they would have to resort to gesturing at things that are gricular and hope that we get the general idea. A few of the more scientifically literate speakers of the gruesome language might tell us that gricular things are those things that either reflect light of a wavelength of around 510 nanometers or are shaped so that the points on its boundary are all some fixed distance from a center point and from a continuous figure. But they wouldn't be giving you a *definition* of "gricular." What they would be doing instead is similar to what we do when we say what wavelength of light green reflects. Gricular is not a word in their language that can be explicitly defined in terms of other words in their language, not even by speakers fully competent in the gruesome language. 1.31

Speakers of the gruesome language classify things in a way that seems odd to us. Just consider all of the things that are grincular. Is this a list of things that objectively belong together? If nothing objectively belongs together with anything else, then their way of classifying objects is not objectively better or worse in this way, although there might be practical reasons to prefer continuing to classify objects in the way that we do instead of changing wholesale to their systems of classification. 1.32

I am going to begin to develop an argument for the conclusion that some things objectively belong together. This argument turns on the idea that certain words are *projectable* while others aren't, and then goes on to claim that a classification system is a bad one if the words it uses to classify objects aren't projectable. 1.33

The key technical term to be explained here is "projectible." The very rough idea is that a projectable term is one that you can justifiably use to describe things that you don't directly perceive. Let's work our way up to a definition of this key technical term. We frequently make generalizations on the basis of a limited number of observations, and if it is irrational for us to do this, the kinds of generalizations we want to make in the course of doing science are irrational as well. Here are some examples. It's snowing outside and you don't want to be cold so you put on a heavy coat rather than just go outside in a t-shirt. Why do you do this? Because you know that thicker clothing is generally warmer than thinner clothing. This is a generalization that you believe, but it is one that you believe on the basis of 1.34

sampling a relatively small collection of clothing. There are probably millions of heavy coats in the world, and probably millions of t-shirts as well. You haven't tried on each heavy coat. And you haven't tried on each t-shirt. Does this shake your confidence that, in general, heavy coats are warmer than t-shirts? I doubt that it does—although you haven't sampled *every* coat and t-shirt, you have tested *enough* to be justified in believing the general claim that heavy coats are warmer than t-shirts.

1.35 Here's a second example. You notice that whenever you drop something heavy on your bare foot, like a laptop computer, it hurts quite a lot. You also notice that whenever you drop lighter things on your bare foot, such as a shoe or a pencil, it hurts a lot less. On the basis of these observations, you conclude that heavier things falling on bare feet tend to hurt more than lighter things falling on bare feet. (You should probably also conclude that you need to wear shoes more often.) You definitely believe this conclusion: were you forced to choose between dropping a feather on your foot or a piano, you know which one you'd choose. But you haven't observed every possible object that could be dropped on your foot in order to assess how painful that experience would be. You are prepared to form a very general belief on the basis of a very limited number of observations, and you take yourself to be perfectly rational in doing this.

1.36 The general pattern of reasoning we seem to be using in these situations is the following. Whenever we encounter a *sufficiently large sample* of things that each have a certain feature F, and each member of this sample also has the feature G, we generalize and believe that *all* things that have F also have G. Philosophers call this kind of reasoning "inductive generalization," and the method of deriving conclusions in this way, "induction." It's hard to see how we can get by in the world if we aren't justified in believing the conclusions of inductive generalizations. How do I know that this bread will nourish me rather than cause me to explode? Because every time I have sampled bread in the past, it has nourished me rather than caused explosions. My sample size is limited but I nonetheless conclude that I'll be ok next time I have a sandwich.

1.37 So far, so good. But now we are ready to consider the following puzzle, due originally to a philosopher by the name of Nelson Goodman, called "the New Riddle of Induction."[10] Once we see how the puzzle works, we'll be able to define "projectable."

---

[10] This puzzle appeared in chapter 3 of Nelson Goodman's book *Fact, Fiction, & Forecast* (Cambridge, MA: Harvard University Press, 1995).

Here we go. Let's think about emeralds. Emeralds have lots of features. 1.38
Some emeralds are pretty. Some emeralds are covered in dirt, and some emeralds will probably never be seen by anyone, at least not during our lifetimes. Some emeralds have been visually inspected by us, or at least will be, during our lifetime. Let's say that something is *grue* if and only if it is *either* both inspected by us during our lifetime and is green, *or* it is blue but not inspected by us during our lifetimes. (As you can probably guess, "grue" is another word in the gruesome language.) Because the idea of grue is a little complicated, it's good to read that definition a couple of times, and then work through some examples. All the emeralds in a jewelry store were visibly inspected by someone during your life and they are green—and so they are also grue.

Here's the new riddle of induction. We haven't seen absolutely all of the 1.39
emeralds in the world. Yet we have seen a large sample of them. And all of the ones that we have seen have been green. So now we are ready to make an inductive generalization: since we have seen a sufficiently large sample of emeralds, and all of these emeralds are green, we infer that all emeralds are green. However, notice this. All of the emeralds we have seen are also grue. So, we've observed the exact same number of grue emeralds as we have green emeralds. So, we should also be willing to make a second inductive generalization: given that all of the emeralds that we have seen are grue, and we have seen a heck of a lot of emeralds, we should be willing to infer that all emeralds are grue too. So, as the result of inspecting the same sample of emeralds, we conclude both that all emeralds are green and that all emeralds are grue.

The problem is that the claim that all emeralds are green is incompatible 1.40
with the claim that all emeralds are grue, since there are emeralds that we have never seen and that won't be observed during our lifetime. Consider an unobserved emerald, which I will call "Eddy." If all emeralds are green and all emeralds are grue, Eddy is both green and grue. But no one has or will observe Eddy in our lifetime. So, Eddy can't be green and observed in our lifetime. Since Eddy is grue, and not green and observed in our lifetime, it follows that Eddy must be blue and not observed in our lifetime. So, Eddy is blue. But just a moment ago we said Eddy is green. Nothing can be both blue and green at the same time. So, the same sample set and the same method of forming more general beliefs from that sample set led to inconsistent results. Not good. That's the puzzle of the new riddle of induction.

What seems plausible then is that we shouldn't indiscriminately make 1.41
inductive inferences. It is ok to infer from our sample size that all emeralds

are green. It is not ok to infer from that same sample that all emeralds are grue, even though every emerald in that sample is both green and grue. Don't believe that whenever we encounter a *sufficiently large sample* of things that each have a certain feature F, and everything that has F also has another feature G, we should generalize and so believe that *all* things that have F also have G. Instead, consider something more complex and careful. Only some of those Fs and Gs are ones that license inductive inferences. We are finally ready to define "projectible": an expression is *projectable* if and only if it stands for a feature that we can justifiably make inductive inferences about. "Green" is plausibly a projectable expression, but if it is, then "grue" can't also be projectable.

1.42 But what makes an expression a projectible expression? One answer to this question is that an expression is projectible to the extent that the expression corresponds to objects that objectively belong together. If this is the correct answer, we also have an argument that some objects objectively belong together. The argument is this. We are sometimes justified in making inductive inferences. But we are sometimes justified in making inductive inferences only if some of the words we use to make those inferences are projectible. And those words are projectible only if they classify objects that objectively belong together. So, some objects objectively belong together.

1.5 Two Questions about Classification

1.43 I hope that we now have a decent enough grip on the distinction between what I have called "subjectively belonging together" and "objectively belong together" that we can proceed to ask some interesting questions about this idea of things objectively belonging together. Let's now turn to two of the most important general questions.

1.44 I want to distinguish the question of *what it takes for some things to objectively belong together* from the question of *which things objectively belong together*. Both questions are important questions, and knowing the answer to one could help us learn the answer to the other. But the answers to these questions are by no means guaranteed to be the same.

1.45 The first question, the one that asks "What does it take for some things to objectively belong together?," is an example of a type of question that philosophers are prone to ask. The type of question of which this first question is an instance is "What does it take for a thing to be X?" where X is some feature that a thing could have. Here are some other examples of this

type of question: "What does it take for a belief to count as knowledge?"; "What does it take for an action to be one that we morally ought to do?"; and "What does it take for an action to be done freely?" Let's briefly think about one of these questions so that we can get a feel for this type of question in general. We'll focus on the first one, the question of what it takes for a belief to count as knowledge.

1.46 There are some beliefs that don't count as knowledge. For example, some people have false beliefs, and no false belief can be knowledge. The other day, I was trying to find my wallet. I had a very strong belief that I had left it on top of the fridge the previous evening. My wife insisted that I hadn't, and I replied that I knew that I had. It turns out that I had left my wallet in my office. Given where my wallet in fact was, I didn't know that I had left my wallet on top of the fridge. I only thought that I knew. Reflecting on this story makes it clear that a belief counts as knowledge only if that belief is true. But it's also clear that merely having a true belief is not sufficient for that belief to count as knowledge. Suppose someone, who we will call "Fred," reads the story I just told and as a consequence comes to believe that my wallet had been in my office. Fred's belief is true, but does it count as knowledge? Well, you might think that Fred actually doesn't have that great of a reason for believing that I left my wallet in my office. It is a well-known fact that philosophy professors make up stories all the time just to provide vivid illustrations of some point that they want to make. Fred doesn't have enough evidence to rule out the possibility that I am just making up a story. So, even though Fred's belief is true, Fred's belief doesn't count as knowledge. It just doesn't have what it takes to count as knowledge. Ok, so what does it take for a belief to count as knowledge? In other words, what are the necessary and sufficient conditions that something has to meet in order for it to be an instance of knowledge? This might be the most important question of epistemology, and if this were an epistemology book, we'd focus on it further.

1.47 The question "What does it take for a belief to count as knowledge?" is a different question from "Which beliefs count as knowledge?" The first question asks for insight into what knowledge is, while the second question asks for a list of beliefs that are known. In principle, someone could answer the second question simply by exhaustively listing all of the beliefs that count as knowledge, which include Jason's knowledge that the sky is blue, Marcy's knowledge that $2 + 5 = 7$, Raul's knowledge that World War II took place prior to the first Gulf war, Shamik's knowledge that emeralds are green, and so on. But this sort of list doesn't by itself answer the question "What

does it take for a belief to count as knowledge?" If we had an answer to *that* question, we'd know *why* all the items on the list of beliefs that count as knowledge belong on this list.

1.48 The same idea applies to the two questions "What does it take for some things to objectively belong to each other?" and "Which groups of things objectively belong together?" An answer to the first question provides necessary and sufficient conditions that a group of things have to meet in order to objectively belong to each other, while an answer to the second question could simply take the form of a list: *these things objectively belong together*, *and so do those other things*, and so on.

1.49 Even though we've distinguished the two questions, as I said earlier, finding the answer to one of these questions might help us find an answer to the other. If we had an accurate list of which groups of things objectively belonged together, we could use that list to test alleged answers to the first question. Suppose, for example, that a proposed answer to the first question said that having X is what it takes for some group of things to objectively belong together, but one of the groups on our accurate list did not have X. We'd be able to deduce that this proposed answer to the first question is mistaken. On the other hand, if we had an answer to the first question, we'd hopefully have some guidance on how to go about making our list of things that do objectively belong together.

1.50 So, we should distinguish our two questions, but keep in mind that the strategies for answering them might not be completely independent of each other. In the following section, we are going to focus on the first question of what it takes for some things to objectively belong to each other.

1.6 Classification and Properties

1.51 Here is a natural train of thought. Things objectively belong together when they are similar in important ways. Things are similar in some way if and only if there is some way that those things all are. We call ways that things are *properties*. So, what it is for things to belong together is for them to share certain important properties.

1.52 But although this train of thought is natural, someone might reasonably worry that we won't travel very far on it. After all, this train of thought leads us to the question "What makes a property an important property?" If our only answer to that question is that important properties are those that things have when they objectively belong together, we are riding the train

on a circular track, and we are no better off than we were when we started. (Circular explanations seem to get you nowhere, intellectually speaking.)

1.53 The question "What makes a property an important property?" matters if we assume both that there are some important properties *and* that there are some unimportant properties as well. But maybe the latter assumption is one that we shouldn't make; maybe we should believe that the only properties that exist are the important properties. If the only properties that exist are the important properties, then we can say that things objectively belong together if and only if they have a property in common. This train of thought is less obviously on a circular track. We still need to face the question of what properties there are though.

1.54 In order to address the question of what properties there are, it might be useful to have a quick discussion about how we attribute properties to things in ordinary language.

1.55 Consider the sentence "Kris is short." The subject of this sentence is the name "Kris." The remainder of the sentence is the predicate "is short." Very roughly, in general, a predicate is that part of a sentence which contains a verb and is used to state something about what is named by the subject of that sentence. Other examples: in the sentence, "Ranger is a dog," "Ranger" is the subject and "is a dog" is the predicate; in the sentence "Ben eats French fries," "Ben" is the subject and "eats French fries" is the predicate; in the sentence "José was an awesome teacher," "José" is the subject and "was an awesome teacher" is the predicate. I hope this rough definition of the word "predicate" and the examples just mentioned have succeeded in giving you the idea of what a predicate is.

1.56 Predicates are used to say something about what is named by the subject of the sentence. Does this mean that, for every predicate, there is a corresponding property? Should we say that, since we can state something true of Ranger when we say "Ranger is a dog," there is a property of being a dog? Ranger is a dog. Is there therefore a property of being a dog that Ranger has? This thought might seem plausible at first—what is the *something* that is true of Ranger if it is not a property had by Ranger? But if there is a property corresponding to every predicate, then there must be unimportant properties as well as important properties. For the following sentence is true: "Ranger is a dog or a fish or an automobile," and so corresponding to the predicate "is a dog or a fish or an automobile" is a property, namely, the property of being a dog or a fish or an automobile. If there is a property of being a dog or a fish or an automobile, it is a property had by Ranger, the fish in my daughter's fish tank, and the leaf-filled convertible rusting on the

street down the block. But these things do not objectively belong together simply because they have this property. The property of being a dog or a fish or an automobile is not an important property.

1.57 So, we have a choice to make. One option is to believe that for every predicate, there is a corresponding property and some, but not all, of these properties are *important*, and then give a theory of what it is for a property to be an important property. Another option is to deny that every predicate corresponds to a property: on this option, although it is, for example, true of both Ranger and Mars that they are either dogs or planets, nonetheless there is no property of being a dog or a planet. On this second option we don't obviously need a theory of when some property is an important property, since we can just say that all properties are important properties. (That said, if being an important property is something that comes in grades or degrees—if all properties are important but some important properties are *more* important than other important properties—then we still need a theory of grades or degrees of importance.) We do need a theory to tell us which predicates have properties corresponding to them. I suspect that on either option, the theories that we will give will look very similar to each other when they are fully developed, but I encourage you to see whether this suspicion ends up being correct. (We will further discuss the idea of an important property in Sections 2.4 and 2.9.)

1.58 Here's another idea worth considering: maybe we can't define or explain what it is for some things to objectively belong together in terms of anything simpler. Instead of trying to define this concept, we could try to use it to define other concepts—such as, for example, the concept of an important property. It's not clear to me how to do this, but that doesn't mean this isn't a project worth pursuing, especially since we already know that some expressions in our language can't be explicitly defined.

1.59 Either way, the metaphysics of classification and the metaphysics of properties are deeply connected. And so, if we want to know more about the former, it will help to know more about the latter. Chapter 2 will focus on the metaphysics of properties, although we will on occasion revisit some of the issues discussed in this chapter.

1.7 Doing Metaphysics

1.60 We've barely scratched the surface of the metaphysics of classification. Here are some further questions to consider:

We distinguished between subjective and objective ways of classifying objects. Are there other ways of classifying objects that aren't neatly categorized as either subjective or objective?

Are there *degrees* of objective belonging? That is, does it make sense to talk about quantities of objective belonging, for example, these objects belong to each other to degree *n*? If not, is there a best way to understand *grades* of objective belonging?

Are there other classifications of people that seem at first to be objective but really are subjective? If so, what are they, and why do they seem (to at least some people) to be objective?

How would you solve the new riddle of induction? Are there good solutions that don't claim that some things objectively belong together?

If not every predicate stands for a property, when does a predicate stand for a property?

Further Reading

In addition to the readings I mentioned earlier in this chapter, here are some suggestions for further reading, along with very brief descriptions of them:

Anjan Chakravartty (2011) "Scientific Realism and Ontological Relativity," *The Monist*, vol. 94, no. 2, pp. 157–180.[11]

This is an important paper on the metaphysics of classification from a perspective from the philosophy of science rather than metaphysics.

Sally Haslanger (2012) *Resisting Reality: Social Construction and Social Critique*, Oxford: Oxford University Press.[12]

This collection of essays contains important discussions of the metaphysics of social groups, including discussions of the metaphysics of race and gender.

Ned Hall's discussion of David Lewis's views on the natural and non-natural property distinction in the *Stanford Encyclopedia of Philosophy*.[13]

An important philosopher discusses one way of cashing out the distinction between important and non-important properties discussed earlier in this chapter. (And in Chapter 2 as well.)

---

[11] http://monist.oxfordjournals.org/content/monist/94/2/157.full.pdf

[12] https://global.oup.com/academic/product/resisting-reality-9780199892624?cc=us&lang=en&#

[13] http://plato.stanford.edu/entries/lewis-metaphysics/natural-distinction.html