

## Levels

John Gibbons

Natural realism is the view that there is a real, metaphysical distinction between those properties, perhaps like *being a quark*, that are natural kinds and those, perhaps like *being either red or round* that are not. This is a fairly respectable doctrine, though by no means universally accepted. Perhaps slightly less well entrenched is the idea that there is a relation of causal relevance that sometimes holds between a property of a cause and a property of an effect. If causation is a relation between tokens, causal relevance is a relation between types. Finally, there's the idea that there are many levels. Sometimes, people will talk about the relation between the mental level and the biological or physical level without meaning too much by this. But the straight-faced suggestion that levels are mind-independent features of reality will occasionally raise a skeptical eyebrow.

It may be that anyone suspicious about any of these three ideas is suspicious of them all. This is as it should be. I don't say this because I think the three characters are maximally suspicious, but because the three form a tightly knit family that to a certain extent stand or fall together. If there are natural kinds, then there are real resemblances, or objective similarities between members of those kinds. But not just any similarity counts. There's at least one respect in which anything three miles from a burning barn is similar to anything else three miles from a burning barn. The similarities that count as far as natural kinds are concerned are causal similarities. Being a member of a natural kind must make a difference to your causal interactions. But the idea of a feature that makes a difference to your causal interactions just is the idea of a causally relevant feature.

If you put the white ball into the pocket of an ordinary pool table in a bar, it will come

back out. Something about the white ball makes it come back out. And that something is probably not its color. It's some other causally relevant feature of the ball. Perhaps the notion of causal relevance can be analyzed in terms of counterfactuals. Perhaps it can be analyzed in terms of the laws of nature. Perhaps it can't be analyzed. Nothing I say in this paper will rely on your analysis or theory about causal relevance. I'll only rely on the intuitive idea of a feature that makes a causal difference. Being a member of a natural kind is the kind of feature that makes a causal difference.

The notion of a level may be the shadiest character of the three because it is commonly associated with nonreductive materialism and Emergentism. According to these views, the biological, geological, psychological, economic, and social levels are both irreducible and autonomous. The higher levels are autonomous in the sense that they contain natural kinds that are causally relevant properties that make a genuine difference to what happens in the world, but these kinds are irreducible in the sense that they are not nomologically coextensive with the natural kinds of physics. Let's just say that nonreductive materialists face certain difficulties in filling out the details of their view. In addition to irreducibility and autonomy, there's also the idea that the lower levels, or at least, the lowest level determines the higher levels. But then we need to know what kind of determination is involved and whether that kind of determination is really compatible with irreducibility and autonomy.

These are serious questions for the nonreductive materialist, and I intend to set them aside.

Whether or not there are many levels is independent of whether or not the levels reduce. Suppose chemistry reduces to physics, and suppose this means that we have not only nomologically necessary biconditional bridge laws but metaphysically necessary property identities. So all chemical kinds are physical kinds. It's consistent with this that not all physical kinds are chemical kinds. It's at least conceivable that there are physical differences that make no chemical

difference. If this is the case, you have the set of physical kinds, only a proper subset of which are chemical kinds. This could be a case in which chemistry forms a level.

The difference between those kinds in the proper subset and those outside is not a difference at the physical level. If you're clever enough, you can, using only physical terminology, produce a disjunctive predicate that will pick out the correct proper subset not only in the actual world but at all possible worlds. But if your predicate has gotten the extension right, then the property it expresses is the kind we also call being a chemical kind. For a strict reading of "chemistry is physics," you need not only a reduction, but the stronger claim that all and only chemical kinds are physical kinds. If biological or psychological kinds do reduce, perhaps in some less straightforward sense, to physics, there may still be biological or psychological similarities inaccessible from the perspective of physics unless we continue to do biology and psychology using terminology that looks highly gerrymandered and disjunctive from the point of view of physics.

The claim that there are many levels does not commit you to any account of the relations between them. Oppenheim and Putnam<sup>[1]</sup> were concerned not only with their celebrated working hypothesis of micro-reduction, a thesis about the relations between levels. They also gave an account of what levels are. The functionalist literature presents a very different account of what levels are. These are the two main alternatives in contemporary debates over the relations between levels, yet neither gives an adequate account. In the first section of this paper, I criticize these two alternative accounts of levels. In the second section, I present and defend my own account according to which levels are higher-order natural kinds. In the final section, I show how levels differ from other natural kinds. For the most part, the psychological level is not given pride of place. It is simply one example among many. If we want to know how the mind fits into the natural world, we need to have some idea of how the natural world fits together.

## *What Aren't They?*

The traditional Oppenheim and Putnam conception of levels is based on the notion of composition. The part-whole relation determines the hierarchy of levels. Except for the lowest level, the objects on each level have a decomposition into things on the next level down. To refresh your memory, I will present their list of levels.

- 6        Social Groups
- 5        Multicellular organisms
- 4        Cells
- 3        Molecules
- 2        Atoms
- 1        Elementary particles

This is a list of objects, as opposed to kinds, laws, or sciences. Where you are on the list is determined by what you are made of. You are on a level either by being an object of the relevant sort or by being composed of objects of that sort. Of course, there is the explicit assumption that everything is composed of elementary particles, so everything is on level 1, but your proper level is the highest level to which you belong.

This account of levels presupposes some notion of a natural kind.<sup>[2]</sup> Everything with a decomposition into parts has, at a time, a northern half and a southern half. “Northern halves of things” does not appear on the list of levels because *being a northern half* is not a natural kind. But natural kinds play another role as well. In addition to objects, we want to order features, or kinds, or properties into levels as well. Perhaps, given our ordering of objects and the notion of a natural kind, we can define an ordering of kinds. For any natural kind property *F*, if objects on a given level have *F*, and no objects on any lower level have *F*, then *being F* is on that level. So if cells and their aggregates have certain characteristic properties not exemplified by molecules, atoms, or particles, then those properties are on level 4.

This picture does have some mild peculiarities. First, planets and tectonic plates are on

level 3 with the molecules. There is nothing in the spirit of the proposal to rule out a branching structure, but if planets and plates do appear at higher levels, they must be individuated by what they are made of. Second, assuming that we're organisms, we are on the same level as plants. If you are concerned with the relation between levels and sciences, psychology and botany are on the same level. I find this mildly peculiar.

The main objection to this view of levels in the age of multiple realization should be obvious. If there are nonbiological thinkers, or thinkers that are not composed of cells, then at least for the purposes of doing psychology, they should be on the same level as us. If these nonbiological thinkers can form social groups, we will have items on level 6 that are not composed of objects on level 5. Similarly, if it is possible that some living organisms are not composed of cells, we will have the same situation with respect to levels 5 and 4. Things are not on levels because of what they're made of. Things are on the level they're on because of their causal powers, and similarity of composition is not necessary for similarity of causal powers.

There's something fishy about an object-based conception of levels, especially when the same object can appear on different levels. Take a marble and a watermelon and throw them out your window. Do they fall at the same rate? Now put the marble and a watermelon on your table and apply the same amount of horizontal force to each. Do they move across the table at the same rate? If you are interested in these sorts of questions, the fact that one is composed of cells and the other isn't makes no difference. And even if no one was ever interested in anything, facts about cellular composition would play no causal role in the production of these effects. Reproduce the relevant features of the marble in cellular material (ivory or wood) and you reproduce the relevant effects.

The two marbles and the watermelon are all on the same level here not because everything is composed of atoms in the void (or quarks in the field or whatever it turns out to

be). They are on the same level because similarities and differences in certain features result in similarities and differences in features of the same sort. It's the features themselves that determine the levels and not the objects that have them. If events, states, and processes can have features of the same sort, they'll be on that level too.

In addition to its clarity, this account of levels does have some admirable qualities. Most importantly, this is a metaphysical, as opposed to epistemic or conceptual account of levels. Where something is in the hierarchy does not depend on anyone's interests, explanatory purposes, theories, concepts or what have you. Science may be the best way to find out which things are on which levels, but we should expect a certain amount of slack between levels and the sciences. Perhaps, on this view, there is too much slack. Not only are botany and psychology on the same level, but biology is as much concerned with level 4 as it is with level 5. But if we think of the levels metaphysically and the sciences as at least partly involving explanatory practices, there should be some room for the two to come apart. Biochemistry and the search for biological bases for mental disorders may cross borders every day. Explanations are epistemic creatures, and they're everywhere. But our epistemic interest in crossing borders doesn't mean the borders aren't out there.

The second common picture of levels is the functionalist account. The functionalist account is based on the notion of a definition or specification. A functional property is a special case of a second-order (or higher-order) property.<sup>[3]</sup> A second-order property is one whose correct definition, or analysis, or specification involves quantification over first-order properties. On one version of the idea, the definition of a particular second-order property A looks like this.

(F) Being A just is having some first-order property B which is C (or which meets condition C).

Two things make functionalism special. On the one hand, the condition involved in "C" is

primarily a causal condition. On the other hand, using simultaneous implicit definition, you can define all of the mental terminology at once. This allows you to take into account interactions between mental states and avoid standard objections to behaviorism.<sup>[4]</sup>

If we want a metaphysical rather than epistemic account of levels, the functionalist account won't do. Being first-order or being second-order is relative to a definition or specification, and these definitions and specifications have more to do with our ways of thinking and talking than with the properties themselves. Suppose that colors are first-order in some sense. Now consider the following two definitions of the second-order property of being G.

(G1) Being G just is having some lower-order property P such that P = being green.

(G2) Being G just is having some property P that stands in the appropriate relations (e.g., is in between on the spectrum) to yellow and blue.

It's clear from the definitions that being G just is being green. Each property is as second-order as the other (as itself). One way of talking about it uses quantification, while the other doesn't.

The functionalist may well suppose that in the cases that matter, e.g., psychology and biology, the functionally defined property will be multiply realized, in which case we cannot identify the higher-order property with the lower-order property.<sup>[5]</sup> But this brings problems of its own. Suppose we think of mental properties along the following lines.

(M1) Being M just is having some lower-order property P that plays causal role R.

It's quite clear from the definition that the lower-order property plays the relevant causal role if anything does. If the mental property is not identical to the lower-order property, then it's difficult to see how the mental property could play the relevant role as well. It's one thing for different properties to play the same role on different occasions, but quite another for two different properties to play the same role on the same occasion, especially if the properties are related in the way suggested by (M1). If somebody shoots Jones and that somebody is Smith, the

fact that somebody shot Smith does not seem to play a causal role over and above, in addition to, or alongside the fact that Smith shot Jones.

If we identify the mental property with the property that plays the role, it is no more higher-order than it is lower-order. If we don't identify the mental property with the property that plays the role, then functionalism entails type epiphenomenalism. Perhaps we should give up on a metaphysical account of levels. If the functionalist who identifies the mental property with the role property can give an account according to which those properties are not epiphenomenal, and if the cost of this view is the acceptance of an epistemic or conceptual account of levels, something that many people are inclined to believe anyway,<sup>[6]</sup> then the view is well worth the price. On this alternative, what gets ordered into levels are conceptions, or descriptions, or explanations, or ways of thinking about. On this view, one and the same thing can appear higher-level from one perspective and lower-level from another. What you look at is no different. The only difference is in how you look at it. Though I think that any conceptualist or non-metaphysical account of levels will result in the epiphenomenality of higher-level kinds, I will restrict my attention to one particular functionalist proposal.

Instead of thinking about mental properties along the lines of (M1), perhaps (M2) provides a better model.

(M2) Being M just is having the property P that plays role R.

Given multiple realization, you should think of the description in (M2) along the lines of the description "the book on the table," and you should think of the latter at least roughly along the lines of "the (contextually relevant) book on the (contextually relevant) table."<sup>[7]</sup> On this picture, "pain," or roughly equivalently, "the property that plays the pain role," refers on those occasions that you're talking about humans (or Jones) to c-fiber firing, if that's what plays the role in humans (or Jones). The very same expression, and importantly, on some way of individuating

concepts, the very same concept or way-of-thinking-about refers or applies to some hydraulic state of the feet on those occasions when you're discussing Martians (or Marvin), if that's what plays the role for them. <sup>[8]</sup>

It's easy to see how this leads to a conceptualist account of levels. If you think of something as pain on a particular occasion, and then you think of something as c-fiber firing, it turns out that you're thinking about the same thing. It's not just that you're thinking about the same token state or event. You're thinking about the same type or kind or property. You're just thinking about that kind in different ways. When you think of it as pain it seems higher level, but this doesn't really tell you anything about the nature of the kind thought about on that occasion. Though I have trouble with this result, the real problem for the view comes when we think about different kinds in the same way.

When you think about your own pain, you're thinking about c-fiber firing. When you think about Marvin's pain, you're thinking about some hydraulic state. Of course, you think about these very different kinds in much the same way. You think of them all as pain. But this doesn't tell you much about the nature of the kinds thought about. Think about the different tokens of those different kinds, and ask yourself what they all really have in common. Well, maybe they're all physical. But is there a distinguishing similarity, i.e., something they have in common which distinguishes them from items to which the concept of pain does not apply? Is there a natural kind picked out by the concept of pain?

To the extent that we like (M2), we should answer this question in the negative. On that picture, your pain plays its causal role in virtue of its neurological features. Marvin's pain plays its role in virtue of its hydraulic features. There doesn't seem to be any causally relevant feature that they have in common. If you're tempted by the idea that what they all have in common is that they play the pain role in virtue of some causally relevant property or another, then you've

retreated away from (M2) and back to (M1). Since the causally relevant property is different in different cases, what all the pains have in common may be a second-order, functionally defined property, but it's not a causal power.

There's currently a fair amount of concern that mental properties and higher-level properties more generally may be epiphenomenal.<sup>[9]</sup> There's also a fairly widespread consensus that higher-level properties, if there are any, are functional properties. These facts are not unconnected. If higher-level properties were functional properties, they would be epiphenomenal. Giving up functionalism means giving up the idea that higher-level kinds can be given second-order definitions, either second-order definitions that involve quantification over first-order kinds, like (M1), or second-order definitions that involve definite descriptions purporting to refer to first-order kinds, like (M2). This is a defining feature of functionalism,<sup>[10]</sup> and this is what leads to epiphenomenalism. The basic idea of the second-order definition, however the details go, is to pick out the higher-level kinds not on the basis of their causal powers, but on the basis of the causal powers of their minions: their first-order realizers, subvenient bases, implementing mechanisms, or what have you. Once you give power to the minions, revolution is inevitable.

If we do give up on functionalism, we at least leave open the possibility that upper-level kinds have their own distinctive causal powers. Making sense of this apparently simple idea turns out to be extraordinarily difficult. We need some account of what it is for a kind to have causal powers, i.e., some account of causal relevance as a relation among types, and we need some account of how the causal powers of an upper-level kind relate to the causal powers of a lower-level kind in a particular that is a member of both kinds. In addition to these extremely difficult questions, it would also be nice to have some idea of what levels are, and what's on them, and what makes it the case that those things are on those levels. It's to these somewhat

easier questions that we now turn in the hope that our answers to them may shed some light on the more difficult questions.

### *What Are They?*

If the composition of objects can't individuate levels, and functional definitions can't either, what else is there? Why, the kinds themselves of course. Levels are just natural kinds of natural kinds. When the same object appears on different levels, it does so because it has different properties or because it falls under different natural kinds. These natural kinds themselves also fall under natural kinds. A falling watermelon is a fruit, and it has seeds. But it also has mass and a particular amount of mass, and it's subject to various gravitational forces. The items on our first list of properties both seem to have something in common. They're both biological kinds. The same goes for the items on the second list. They're all physical kinds ("physical" in the sense of physics, not in the sense of not immaterial). If *being a biological kind* is itself a natural kind, it determines a level.

Let's start with some terminology. In addition to its use by functionalists, "second-order" also has a more metaphysical sense. On this way of talking, there are properties of and relations among particulars, and these are first-order properties and relations. Then there are properties of and relations among first-order properties and relations, and these are second-order properties and relations. Rather than introducing new terminology, I'll use "second-order" in this sense. If there is a risk of confusion, I'll use "second-order in the functionalist's sense" to talk about definitions that involve quantification. So *being a color* is a second-order property in both senses, and *having a color*, though second-order in the functionalist's sense, is a first-order property in the metaphysical sense: it's a property of particulars. I'm also going to assume that the kind *penguin* just is the property of *being a penguin*. So all kinds are properties, but not all properties are natural kinds. <sup>[11]</sup> With this in place, we can say that levels are higher-order natural kinds.

When I talk about natural kinds, I don't mean anything about hidden essences. It's conceivable that there are creatures whose sensory apparatus is so fine-tuned that it was just obvious to them that water is H<sub>2</sub>O. And what it takes to be a zither is as hidden from me as what it takes to be plutonium. Also, it is very unlikely that every natural kind term expresses an essential property of the things that fall under the term. Having a particular amount of mass may well be a natural kind. But if you gain an ounce, you will not pop out of existence. So what do I mean when I talk about natural kinds?

On the standard picture of natural kinds, looking for what all and only members of a natural kind have in common is a perfectly legitimate, though, of course, straightforwardly empirical, scientific procedure. This is where the familiar property identity statements are most at home.<sup>[12]</sup> Being water just is being H<sub>2</sub>O. Being gold is having atomic number 79. Perhaps, being a raven is being a member of a certain reproductively isolated population. These identities provide necessary and sufficient conditions for membership in the kind. But necessary and sufficient conditions don't eliminate vagueness. They transmit it. Just as the vagueness of "bachelor" is inherited from the vagueness of "man," the vagueness of "raven" is inherited from the vagueness of the boundaries of the relevant population.

On this picture, it's a natural fact that water is H<sub>2</sub>O. This is not made true by convention, or definition, or stipulation, or meaning. It may be a conventional fact that "water" means *water*, but nothing we've ever done has made it the case, and nothing we could ever do will make it not the case that water is H<sub>2</sub>O. But there isn't just one natural fact about water. There are all kinds of natural facts about water: facts about its boiling and freezing points, facts about its thermal and electrical conductivity, and perhaps facts about its transparency. None of these facts are a priori. They are not true in virtue of conventions about meaning; they're true in virtue of the laws of nature. We know these facts in the same way we know the laws: on the basis of experience.

The kind *water* is connected, not by convention but as a matter of causal law, with two very different sets of properties. First, there's *being H<sub>2</sub>O*, and this is what makes it the case that something is water. Second, there's the (usually significantly larger) collection or cluster of kinds to which water is lawfully related, boiling point and all the rest. The idea that natural kinds involve real resemblances means that these relations among kinds are neither accidental nor conventional or conceptual. Though we might discover the similarities between different samples of water, the similarities are not imposed by the mind. But then, it might not really make any sense to talk about the mind imposing similarities, so perhaps the idea is this. Though it's not just an accident that bachelors are unmarried, this kind of nonaccidentality is conceptual rather than causal. If it's not just an accident that water is a liquid at room temperature, this nonaccidentality is causal rather than conceptual.

This is the basic idea behind natural kinds. Perhaps thinking of natural kinds as homeostatic property clusters provides a better picture of the facts or a better account of what natural kinds really are.<sup>[13]</sup> But I'm only relying on the basic idea, and the basic idea seems to have (at least) the following two epistemic consequences. It's not a priori whether a purported natural kind term succeeds in picking out a kind. But if a term does pick out a natural kind, then induction involving that term, or at least, a certain kind of induction, is reliable. If reliability has any epistemic significance at all, this is an important consequence. If similarities between different samples of gold are guaranteed by the causal structure of the world, whatever exactly that means, and one of the similarities so guaranteed is a similarity in melting point, then believing something about the melting point of samples of gold in general on the basis of examination of instances is likely to get things right.

I think the notion of a level is just beneath the surface of the notion of a natural kind. Think about the cluster of properties lawfully connected to water, and compare it with the cluster of

properties connected to gold. These are clearly different clusters or collections: different melting points, different boiling points, and differences in thermal and electrical conductivity. But the similarities between the kinds jump out at you when you compare those two clusters with the cluster of kinds connected to ravens, or tectonic plates, or planets, or, neutrinos. Water and gold may have different melting points, but at least they have melting points. In this, they differ from all the other kinds on our list. But it's not just gold and water that have melting points, so do copper and salt. In fact there are quite a number of different kinds that have melting points. But it's not just that water, gold, copper, salt and their friends have melting points. They have differing degrees of thermal and electrical conductivity, malleability, ductility, and all the rest. The degrees differ, but they differ across kinds rather than within kinds.

In addition to the differences in the clusters, there are also important differences between what it takes to be gold and what it takes to be water. But now that you're looking for it, the similarity is obvious. Atomic structure is what makes any particular chemical kind the kind that it is. It's not an accident that similarities in atomic structure also explain similarities in melting points, conductivities, and all the rest. If these similarities between the chemical kinds, not just water, gold, copper, and salt, but also melting points, conductivities, and malleability, are as real as the similarities between different samples of gold, then *being chemical* is a natural kind, and chemistry forms a level. To say that the similarities are as real as the similarities between samples of gold is to say, approximately, that the complex relations that hold among these kinds is guaranteed by the causal laws. They're not accidental, and they're not conceptual either. The similarities among particulars are determined not by convention but by the causal laws, and this is what groups them into kinds. The laws also connect different kinds together: water to its boiling point and gold to its boiling point. And similarities among the laws (gold, water, copper, and salt are all lawfully connected to melting points, conductivities, and atomic structures)

determine a level. As we'll see, we should think that levels are higher-order kinds because they play pretty much the same role in stating and confirming laws that other natural kinds do. Given this reliance on laws, perhaps I should say something about them.

Statements of law do not involve first-order quantification over particulars. They express propositions about second-order relations among universals. I take it that the case for this has already been made.<sup>[14]</sup> In this respect statements of law resemble (F).

(F) The French eat frog's legs.

The truth of (F) does not require every French person everywhere to eat frog's legs, nor is it sufficient for one of them to do so once. Since the sentence never means either of these things, it's certainly not ambiguous between the two. To make sense of (F), we don't need to look for a third quantifier over individuals, e.g., "most French people." Even if only 49% of French people ate frog's legs, (F) would still be true. Nor do we need to start with a universal conditional and then somehow weaken it with a *ceteris paribus* clause. We don't need to find or modify a quantifier over individuals because (F) is about a relation between universals. And understanding (F) is knowing what that relation is. Perhaps, even in (F), there is a suggestion of nonaccidentality: culinary habits are perpetuated within a culture. But our concern is not with (F). It's with statements of law.

If the relation between the universals referred to in a statement of law is not understood in terms of an exceptionless correlation or a first-order, universal conditional, how should we understand it? When it comes to laws, nonaccidentality is more important than universality. If it's a law that Fs are G, then it's not an accident that Fs are G. Something makes Fs G. But it's not different things on different occasions. It's the same thing that makes different Fs G. And the thing that makes different Fs G is the property of being F. To use the jargon, the relation between universals that makes true law statements true is the relation of causal relevance. Just as (F)

neither is nor entails a universal conditional, neither are nor do laws. The question of lawhood is not determined by how often or how rarely things go wrong, e.g., the percentage of Fs that are not G. The relation between the properties when things go right makes something a law, e.g., that it's the property of being F that makes these things G.

The notion of a natural kind is best understood in terms of its relations to the notions of a law, causal relevance, induction, and explanation. Perhaps the best way to argue for the claim that levels are natural kinds of natural kinds is to show that the notion of a level stands in basically the same relations to that set of notions. Let's start with higher-order kinds in general. Various relations among properties, such as inclusion, determination, and parenthood, determine the structure of Plato's heaven. The parenthood relation between properties or kinds is a slight modification of a relation defined by Goodman. The nominalist version goes like this. A predicate "P" is a parent of predicate "Q" if among the classes "P" applies to is the extension of "Q."<sup>[15]</sup> So think of the extension of "P" as a set of sets. The extension of "Q" is a member of this set, not a proper subset. This is how the parenthood relation differs from the determinable-determinate relation. The set of scarlet things is a proper subset of the set of red things but a member of the set determined by "shade of red." Without altering Goodman's theory too much, we can add a little harmless realism about universals and Natural Realism. *Being a shade of red* isn't really a set of sets. It's a kind of kinds, and scarlet is one of those kinds.

*Having seven grams of mass* is a determinate of *having mass* and a child of *amount of mass*. The property of *having some mass or another* is a first-order determinable property, and at least in certain cases, there may be some reason to doubt its status as a causal power. If something produces an effect of a certain sort partly in virtue of having seven grams of mass, if its having seven grams of mass is causally relevant on this particular occasion, then it's hard to see how that same thing's having some mass or another could play a causal role, or how it could play the

causal role already played by one of its determinates. Perhaps, if all things with mass behave the same way in a certain respect, then the determinable property might be causally relevant with respect to this feature of the effect. But if you'll act the same way no matter how much mass you have, it's hard to see how having seven grams of mass makes a difference over and above having some mass.<sup>[16]</sup> In short, with respect to any particular feature of the effect, determinables and determinates compete for causal power.

The property of *being an amount of mass* is a fully determinate, second-order property. But the fact that it's a property of properties may seem to cause trouble for its being a causal power. *Having seven grams of mass* might be a causal power or a causally relevant feature even though the property itself doesn't cause things. At least things that have the property cause things. But the things that have second-order properties like *being an amount of mass* don't cause things, so how can the higher-order property be a causal power? I don't think that what makes something a causal power are facts about which particulars are connected in the right way to which other particulars. What makes particulars causally interact are facts about what properties they have and facts about relations among those properties. Though the properties themselves never cause anything, lawlike relations among properties can only be described as causal. Causal relevance, a relation between properties, is a kind of causal relation. Let me explain.

The law that  $f=ma$  says something about any amount of mass  $m$ . We use parents to quantify over the children. The particular amounts of mass, the children, stand in interesting relations to each other which can be mimicked by relations between numbers. This is why you can replace the letters with numerals, do some math, and find out about the world. Any particulars with the same mass are indistinguishable with respect to this law. What goes for one goes for the other. In one respect, the children are each treated differently. Given the same force, more mass means less

acceleration. But in terms of their higher-order relations to amounts of force and acceleration, the children are indistinguishable with respect to this law as well. Where you are in the force family varies in the same way with your position in the mass and acceleration families, no matter where you are in any of these. This, at least, is what the law says.

It's not an accident that particulars with seven grams of mass behave in similar ways in similar circumstances. The causally relevant feature they share is what makes them behave in that way. It's also not an accident that different amounts of mass confer similar sorts of causal powers. The feature that the features share, being an amount of mass, makes them confer the causal powers that they do. This is why we use the parents to state the laws. If second-order properties can figure in laws and explain similarities in causal powers, I see no reason not to think of them as causal powers. If it's not just an accident that water, gold, copper, salt and their friends have melting points, degrees of thermal and electrical conductivity, malleability, ductility, and all the rest, it may be that what these kinds have in common, the property of *being chemical*, makes them similar in these respects.

We don't just need higher-order kinds to state the laws; we need them to confirm the laws as well. To use a familiar example, black ravens confirm the claim that all ravens are black, but blue cheese does not confirm, at least in the same way, the claim that all non-black things are non-ravens.<sup>[17]</sup> The difference is that while *being a raven* is a natural kind, *being non-black* and *being a non-raven* are not natural kinds. Presumably, to finish the story, we should say that *being black* is also a natural kind, at least, an ornithological kind.

I don't know if *being black* is an ornithological kind, but I think I know what would count as evidence. Just as kinds have parents, so do claims. "Fs are G" is the parent of "As are B" iff "F" expresses a parent of the property expressed by "A" and the same goes for "G" and "B."<sup>[18]</sup> The parent of (R) is (PR).

(R) Ravens are black.

(PR) Kinds of birds are uniform in color.

Evidence that (R) is a law cannot be neatly disentangled from the evidence that (PR) is a law. If you like, they face the tribunal of experience together. This means that a uniform coloring among swans can play an indirect role in the confirmation of (R).

One reason to believe that *being black* is an ornithological kind is the belief that the parent of our raven hypothesis is a law. It's not that each bird is uniformly colored but that each bird has the same kind of coloring as its conspecifics. Reason to doubt this is reason to doubt that color is a kind. The mere fact that we don't individuate species in terms of their colors does not show that color is not a kind. We might individuate gold in terms of its atomic structure rather than any of its observable characteristics but still think that there's a law connecting gold with its melting point.

Of course, it might just be an accident that ravens are black. It might even be an accident that kinds of birds are uniform in color. If we had a biological explanation, any kind of biological explanation for uniform coloring, that would be even better evidence. While it doesn't matter what kind of biological explanation is involved as far as evidence for nonaccidentality is concerned, it does matter that the explanation is biological. There may be some sort of explanation at the level of physics of why each member of a particular set has the color that it does. This may, in some sense, count as an explanation of why those things have the same color if it turns out that they do. But if there's no one thing that makes them have that color, if it's different things in different cases, then we don't have a law. And if there is one thing that makes them have that color but it's not a biological kind, then we may have a law of physics, but we don't have a biological law.

I don't know how the empirical facts turn out in this particular case, and I don't think very much

turns on it. If it turns out for example, that color influences a bird's interactions with other creatures with vision, e.g., potential mates, or predators, or prey, and this explains why the trait is passed on to further generations, this should not come as too much of a shock. The philosophically significant question is not whether color is a biological kind. The questions are what it means to say that it is and what would count as evidence one way or the other. We start off wondering whether it's a law that ravens are black. This leads to the question of whether kinds of birds are uniformly colored and if so, whether it's a law that they are. And this leads to the question of whether certain explanations are biological rather than physical or chemical. The more specific questions have to be answered against the background of the more general questions. This means that there's more to find out. But it also means that more evidence is available. But none of this makes any sense without taking seriously the question of whether *being black*, or *being a raven*, or *being a species of bird*, or *being biological* are natural kinds. If you're trying to confirm the law that ravens are black, it's useful to know that *being a raven* and *being black* are natural kinds. But this knowledge isn't all that useful unless you also know that they're on the same level. Suppose you see some ravens near some gold. The next day, you see some different ravens near some different gold. Knowing what we know about ravens and gold, we're much more likely to think of this as an accident than as a legitimate basis for induction. Why is that? Though "gold" and "raven" are both projectable, they're not, as Davidson puts it, "made for each other."<sup>[19]</sup> It's an empirical question which kinds are made for each other. And it's nature, not us, that makes kinds for each other. And all this talk about kinds being made for each other is really just talk about kinds being on the same level. If *being gold* and *being a raven* are not on the same level, then the relation between them is not the relation of causal relevance; statements about their connections do not express laws; and generalizations about the pair are not confirmable on the basis of induction.

### *What Makes Them Different?*

Levels are natural kinds. This is the most fundamental fact about them. But not just any natural kind is a level. Levels are higher-order natural kinds, but then, so is *being an amount of mass*. So what is distinctive about levels? A level is the proper object of study for a science. This is why “physics,” “chemistry,” “biology,” and the rest can be used to talk about an activity of inquiry or the level that is the object of that inquiry. This epistemic fact about levels is a contingent, relational fact about them. But it does help to fix ideas. If you’re looking for something it makes sense to study, one thing you look for is a kind of unity in the subject matter. If there is unity in the subject matter, and your theory of the subject matter is by and large correct, there will be a kind of unity to your theory. But this theoretical unity is parasitic on the more objective, metaphysical unity of the level.

What kind of objective unity are we talking about? I think the relevant kind of unity is a kind of causal unity. Parenthood is an intralevel affair. All amounts of mass are on the same level. But if  $f=ma$ , the force and acceleration families are on the same level as well. So think of the levels as grandparents. The natural kind, physical kind, has children: mass, charge, charm and all the rest. It also has grandchildren: particular amounts of mass, positive charge, and negative charge. The middle generation kinds, mass and charge, are on the same level because similarities and differences in certain features result in similarities and differences in features of the same sort. The notion of a level or a physical kind gives content to this idea of features of the same sort. Mass and charge influence, in their own ways, position, velocity, momentum and so on, i.e., and other physical kinds, magnitudes, and relations. What this comes down to, in the end, is that what ties the families together into a level is the notion of causal relevance.

Just as it’s not an a priori matter what the natural kinds are, it’s not a priori what levels there are. Reason to believe that physics forms a level comes from physics. Reasons for and

against believing that biology, geology, and psychology form levels come from those sciences as well. But it's also not an a priori matter which kinds are, for example, geological kinds. If there is a family of kinds, members of which are systematically responsible for similarities and differences in geological features, then that family is a family of geological kinds. This is not to suggest that the geological level is causally closed. What Davidson said about the mental holds for the higher levels in general. Too much happens to affect the geological that is not itself a systematic part of the geological. This is why geological laws are not exceptionless. But if something is a systematic part of the geological, that's reason enough to think that it is geological. And if there are enough families of geological properties systematically connected to each other, then exceptions or not, there may be enough unity to the subject matter for it to make sense to study geology.

Does psychology form a level? This is, of course, an empirical question, but the overwhelming, everyday success of folk psychology strongly suggests that it does. Are there families of psychological properties? If we take the evidence of folk psychology at face value, believing that snow is white is a different kind of mental state than believing that grass is green. They make you do different things, or they're causally responsible for different kinds of intentional action. But if there are any psychological laws that quantify over contents,<sup>[20]</sup> there are certain respects in which all beliefs are on a par. In addition to the belief family, where different beliefs with different contents are the children and *being a kind of belief* is the parent, there are also the desire, appearance, intention, and apparently, the intentional action families. If what you know and perceive makes a systematic, causal contribution to how things seem, or what you want, or think, or do on purpose, then these factive states form psychological families of kinds as well. A priori preconceptions about privacy or intrinsicness are simply irrelevant.

At its most basic, the idea that mental properties are causally relevant is the idea that what

we know or think or want is systematically, causally responsible for what we do on purpose. If this is not true, we have no control over our own lives. We are, as Wittgenstein once imagined, like leaves being blown about in the wind all the while thinking to themselves, “Now I’ll go this way ... now I’ll go that way.”<sup>[21]</sup> We can’t control ourselves without higher-level causation. If it turns out that the upper levels do not reduce to the lowest level, then we need irreducible higher-level causation. There are many philosophical difficulties with the idea that there is irreducible, higher-level causation. But the idea that you have to take levels seriously is not one of them.

If you’d like to believe in levels, but you’re a little embarrassed about the whole thing, don’t start by worrying that the mental might reduce to the physical. Start by worrying about non-black non-ravens, grue emeralds, and all that work for a theory of universals that you just can’t do without natural kinds. Once you’ve got natural kinds, you’re in a position to take levels seriously. If you can use them to make sense of nonreductive materialism and the causal relevance of the mental, so much the better.

---

<sup>[1]</sup> In “Unity of Science as a Working Hypothesis,” *Minnesota Studies in the Philosophy of Science*, vol. 2 (1958).

<sup>[2]</sup> Oppenheim and Putnam, p. 9.

<sup>[3]</sup> See, for instance, Ned Block, “Can the Mind Change the World?” reprinted in C. and G. Macdonald (eds.), *Philosophy of Psychology* (Oxford: Basil Blackwell, 1995).

<sup>[4]</sup> See David Lewis, “Psychophysical and Theoretical Identifications,” *The Australasian Journal of Philosophy* 50 (1972) pp. 249-258.

<sup>[5]</sup> For multiple realization, see Hilary Putnam, “The Nature of Mental States” reprinted in *Philosophy of Mind* ed. Davis Chalmers, (Oxford: Oxford University Press, 2002).

<sup>[6]</sup> See, e.g., C.B. Martin and John Heil, “The Ontological Turn,” *Midwest Studies in Philosophy* 23 (1999) pp. 34-60.

<sup>[7]</sup> See David Lewis, “Mad Pain and Martian Pain,” in *Readings in the Philosophy of Psychology* vol. 1, ed. Ned Block, (Cambridge, Mass.: Harvard University Press, 1980), pp. 216-222.

<sup>[8]</sup> Why not let “the property that plays the pain role” refer to the mental property of *pain* on all occasions? Even if there is nothing else wrong with this idea, that is, even if the relevant psychological theory is metaphysically necessary, and even if that theory, with all of its psychological vocabulary replaced with variables still contains

enough information to uniquely specify the set of psychological properties, the resulting picture does not generate an account of levels. If we identify the mental property with the property that plays the role, it is no more higher-order than it is lower-order.

[9] For some of this concern, see Jaegwon Kim, *Mind in a Physical World* (Cambridge, Mass.: MIT Press, 1998).

[10] This is definitive of one thing commonly called “functionalism.” In addition to Ramsey-Lewis, or Causal Functionalism, there’s also the idea that what’s distinctive about biological or psychological kinds can be understood in terms of the notion of a function or purpose. See, for example, Ruth Millikan, *Language, Thought, and Other Biological Categories* (Cambridge, Mass.: MIT Press, 1984) or Elliot Sober, “Putting the Function Back into Functionalism,” in *Mind and Cognition* 2<sup>nd</sup> edition, ed. William Lycan (Oxford: Basil Blackwell, 1999). This idea is not usually intended as a general account of levels. But William Lycan does attempt to give a general account of levels in terms of teleology. See *Consciousness* (Cambridge, Mass.: MIT press, 1987) chapters 4 and 5. The main difficulty for this view of levels is that there doesn’t seem to be anything like a function or purpose on the geological, chemical, or physical levels. Volcanoes may be caused by subterranean pressure. But releasing that pressure is not what they’re for.

[11] Some philosophers claim that there’s no such property as being either red or round. I say there is, but it’s not a natural kind. This is at least partly a terminological dispute. Other philosophers claim that there’s really no distinction between natural and unnatural kinds. This is a far more substantive disagreement, and philosophers of the first kind are on my side. For a sparse conception of properties, see D.M. Armstrong, *A Theory of Universals* (Cambridge: Cambridge University Press, 1978). For reasons to believe in more properties than this while retaining a distinction between natural and unnatural kinds, see David Lewis, “New Work For a Theory of Universals,” *Australasian Journal of Philosophy* 61 (1983).

[12] For this picture of natural kinds, see Saul Kripke, *Naming and Necessity* (Cambridge, Mass.: Harvard University Press, 1980), and Hilary Putnam, “The Meaning of ‘Meaning’,” in *Mind, Language, and Reality*, (Cambridge: Cambridge University Press, 1975), pp. 215-271.

[13] For homeostatic property clusters, see Richard Boyd, “How to be a Moral Realist,” in *Moral Realism* ed. Geoffrey Sayre-McCord, (Ithaca: Cornell University Press, 1988); “What Realism Implies and What it Does Not,” *Dialectica* 43 (1989) pp. 6-29; and “Realism, Anti-Foundationalism, and the Enthusiasm for Natural Kinds,” *Philosophical Studies* 61 (1991) pp. 127-148. Though he may think of the standard picture with necessary and sufficient conditions for kind membership as a rival to his view, my presentation of the basic idea of natural kinds is strongly influenced by Boyd’s work.

[14] See Fred Dretske, “Laws of Nature,” *Philosophy of Science* 44 (1977) pp. 248-268; David Armstrong, *What is a Law of Nature?* (Cambridge, UK: Cambridge University Press, 1983); and Michael Tooley, *Causation: A Realist Approach* (Oxford: Oxford University Press, 1987).

[15] Nelson Goodman, *Fact, Fiction, and Forecast* (Indianapolis: Bobbs-Merrill, 1965) p. 104.

[16] As I see it, these are the intuitions that underlie the notion of proportionality. See Steve Yablo,

[17] Carl Hempel, *Aspects of Scientific Explanation* (New York: The Free Press, 1965).

[18] This is similar to what Goodman, p. 110, calls “overhypotheses.”

[19] Donald Davidson, “Mental Events,” in *Essays on Actions and Events* (Oxford: Oxford University Press, 1980) p. 218.

[20] Perhaps this is a psychological law: If you believe that p on the basis of your belief that q, and you find out that q is false, then you’ll continue to believe that p and make up reasons for so believing if asked. If this is a law, it’s a law that quantifies over contents.

[21] G.E.M. Anscombe attributes the image to “some notes on a lecture of Wittgenstein.” See *Intention* (Oxford: Basil Blackwell, 1958) p. 6.