Antireductionism Has Outgrown Levels

Angela Potochnik

Forthcoming in Alastair Wilson and Katie Robertson (Eds.), *Levels of Explanation*, Oxford University Press.

Abstract

Positing levels of explanation has played an important role in philosophy of science. This facilitated the advocacy of antireductionism of explanations, which, at its most basic, is the idea that scientific explanations citing large (i.e. non-microphysical) entities will persist. The idea that explanations come in levels captures important features of explanatory practices, and it also does well at helping to define different positions one might take regarding explanatory reductionism or antireductionism. Yet the idea that explanations come in levels has also led philosophers astray. This systematically misconstrues the relationship different explanations bear to one another, suggests candidate explanations are less numerous than they in fact are, and occludes recognition of how the selection of explanations can vary across research projects. Antireductionists about explanation should move on from talk of levels. Or so I will argue.

"Entities at different levels;" "explanations in higher-level terms;" "the fundamental level;" "higher-level sciences." These and many similar turns of phrase are used throughout philosophy of science and metaphysics, typically without much in the way of explication. These are used as starting points for discussions rather than the endpoints of argument. By what right are such turns of phrase used? Best I can tell, the rationale for such levels talk is taken to stem from (variously) the mere fact of productive scientific inquiry addressing objects larger than fundamental particles; that scientists of various stripes invoke levels on a regular basis; that some scientific investigations target the very smallest happenings in our world, happenings that seem bound up in one way or another with everything that goes on in our world; that philosophers, scientists, and laypeople alike often have whole discourses without making reference to these smallest goings-on; that there are sometimes multiple candidate explanations of a single explanandum, some of which feature larger entities and their properties than others.

I have done the same. In a 2010 paper on levels of explanation, I simply say, "In general, a lower-level explanation cites properties of objects that stand in a part-whole relationship to objects referenced in the competing higher-level explanation" (Potochnik, 2010, p.64), and I reference lots of important philosophers who talked about scientific explanations in this way. Since then, I have started to examine this assumption that explanations come in levels more carefully and to attend to others who are also questioning this. I have been startled at how little weight these turns of phrase and the assumption behind them can actually bear. And yet, the assumption that explanations come in levels persists as an unexamined starting point of philosophical treatments of explanation. Levels of explanation receive plenty of discussion, but the discussion largely consists in whether there are higher-level explanations and, if so, what relationship they bear to what we know, or might someday know, about the smallest, microphysical happenings in our world. As far as I can tell, that scientific explanations and the entities featured in them are arranged in levels mostly still goes uncontested.

In this paper, I will argue that it is a mistake to invoke levels in discussions of scientific explanation. The invocation of levels played a very important role historically in philosophy of science, as a way to motivate an antireductionist stance about scientific explanation. But our scientific and philosophical understanding has progressed mightily since then, and we can do antireductionism better. It is thus time for philosophers of science to abandon the levels framework in our discussions of scientific explanation.

In Section 1, I outline the role invocation of levels has played in philosophy of science, focusing especially on how they have been used to motivate antireductionism about scientific explanations. In Section 2, I argue that framing antireductionism about scientific explanation as a thesis about levels of explanation has led to problematic commitments—that candidate explanations form a linear or at least partial hierarchy, can be ordered by generality, and bear straightforward metaphysical relationships to one another. In Section 3, I use the difficulties of the levels framing to show how antireductionism can be done better without levels. This involves reconsidering the relationships different explanations bear to one another, recognizing a wider variety of candidate explanations, and appreciating how considerations guiding the selection of explanations can vary across research projects. Finally, in Section 4, I conclude by offering a new 'working hypothesis' about the nature of our scientific explanations: They are many and varied, often featuring large-scale, distant, and structural factors. The decision of explanatory quality is not about how fine-grained our characterization of local factors should be but rather which factors at what scales we should attend to. Reductionism has failed, but so too has the framework of explanatory levels. The levels framing is no longer necessary nor helpful in motivating antireductionism about scientific explanation.

1. Levels in Antireductionism

There is tradition in philosophy, as well as in at least some fields of science, to invoke levels on both sides of debates about reductionism. In philosophy, this tradition traces back at least to Oppenheim and Putnam's influential motivation for the unity of science understood as reduction to physics.[1] Oppenheim

[1] Hempel and Oppenheim (1948) also consider levels of explanation, but their levels of explanation are not compositionally defined but defined in terms of abstractness: "higher levels [of explanation] require the use of more or less abstract theoretical constructs which function in the context of some comprehensive theory" (147). Their illustration is explaining a planet's position with reference to Kepler's laws (lower level) or instead from the general law of gravitation and laws of motion (higher level). Indeed, a strategy of high-level explanation on which Hempel and Oppenheim focus is "explaining a class of phenomena by means of a theory concerning their micro-structure," so the usage of 'levels' is very different from the later Oppenheim and Putnam paper. Thus, although the connection between high-level explanation and greater abstractness that has been influential is established in this paper, the

and Putnam's levels are, in descending order: social groups, multicellular living things, cells, molecules, atoms, and elementary particles. The relation among entities at different levels is one of part-whole composition:

> Any whole which possesses a decomposition into parts all of which are on a given level, will be counted as also belonging to that level. Thus each level includes all higher levels. However, the highest level to which a thing belongs will be considered the "proper" level of that thing" (1958, 9,10).

Note that this is not an endpoint of their analysis but rather the starting point. That is, Oppenheim and Putnam presume this is how our world is ordered—as wholes entirely decomposable into parts occupied at a lower-level—and then investigate what relation we should expect among the fields of science that investigate these levels. They predict, based on empirical evidence of how science seemed to them to be proceeding, that all science would eventually be reduced to microphysics—loosely put, that our best scientific laws would be vindicated and analyzable in terms of microphysical laws.[2]

A decade and a half later, Fodor (1974) responds directly to Oppenheim and Putnam's 'working hypothesis' of the unity of science with the opposed hypothesis of disunity, i.e., the failure to reduce explanations or theories to physical theory. His argument rests on the observation that there is not often a neat relationship between kinds invoked in higher-level explanations and the physical kinds upon which they depend: "interesting generalizations can often be made about events whose physical descriptions have nothing in common" (103). Thus enters the influential idea of multiple realization, and with it the presumption that high-level properties are realized by lower-level properties. We can also credit Fodor with the diagram shown in Figure 1, variants of which have proliferated ever since in

relationship of levels to explanatory reductionism is reversed from what is customary in later philosophical discussions.

[2] For discussion of a very different and largely neglected tradition of the unity of science tracing back to the Vienna Circle, see Potochnik (2011).

discussions of the significance of realization and multiple realization for explanation and causation. Around the same time, Putnam (1975) emphasized the value of the generality of higher-level explanations compared to lower-level explanations, deploying the now-classic example of explaining why a square peg failed to go through a round hole with the same diameter. And Garfinkel (1981) argued that explanation "seeks its own level," that the factors that truly made a difference to some occurrence are found at the same level. For Garfinkel, reductive explanations thus suggested sensitivity

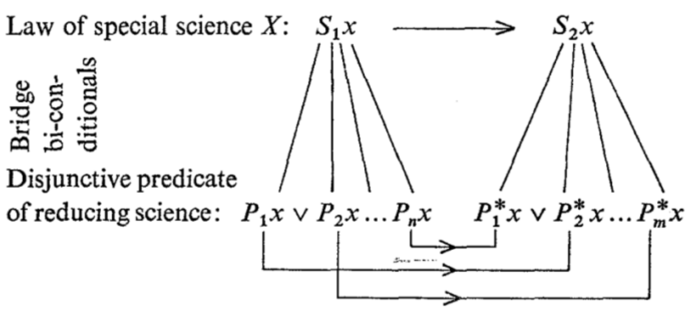

to details that were in fact irrelevant.

Figure 1: Fodor's illustration of why reduction of scientific explanations is unlikely to come to pass.

The idea that higher-level explanations are more general or more abstract than lower-level explanations has since become very influential. Often the justification given is in terms of multiple realization, as in Fodor's influential argument and diagram. Yablo (1992) employs this idea in his proportionality argument for mental causation. Sober (1999) employs the framework of multiple realization giving rise to different levels of explanation to support a pluralism about explanatory strategies, including lower- and higher-level explanations that are, respectively, more specific and more general. Jackson and Pettit (1992) deploy a different approach from Sober's to defend a pluralism that admits both more general higher-level explanations and more specific lower-level explanations. Hauge (2010) and Clarke (2015) each analyzes what specific variety of abstractness might be at play in distinguishing high-level from lower-level explanations.

This combination of ideas has become a general setup for antireductionism about scientific explanation: Candidate explanations come in levels; entities that feature in lower-level explanations compose the entities that feature in higher-level explanations; properties cited in lower-level explanations determine and multiply realize the properties cited in higher-level explanations; higher-level explanations are more general than lower-level explanations. Arguments in favor of mental causation (e.g. Yablo, 1992) and of metaphysical emergence (e.g. Wilson, 2013) defend not just the explanatory but the causal autonomy of higher-level properties conceived in this way. Mechanistic accounts of explanation have a conception of mechanisms consistent with this general setup and deploy it as a competing view to explanatory reductionism (e.g. Craver, 2007).

Discussions of complexity as a bulwark against reductionism also presume this general setup. Here is Herbert Simon, in his classic discussion of complexity and systems theory:

> The central theme that runs through my remarks is that complexity frequently takes the form of hierarchy, and that hierarchic systems have some common properties that are independent of their specific content. Hierarchy, I shall argue, is one of the central structural schemes that the architect of complexity uses.
>
> By a hierarchic system, or hierarchy, I mean a system that is composed of interrelated subsystems, each of the latter being, in turn, hierarchic in structure until we reach some lowest level of elementary subsystem (1962, 468).

William Wimsatt propounded this style of view of levels in philosophy of science, famously describing levels as “local maxima of regularity and predictability” (1972; 2007). This version of antireductionism has differences from the philosophical tradition I surveyed above, but it also resonated with that tradition and has regularly been treated as an allied position. Wimsatt has his own diagram, one that has been more influential in some circles than Fodor’s illustration of multiple realization; see Figure 2.

And, so the idea that explanations come in levels, as perhaps also do causes and organizational relationships, has become entrenched as a key assumption of antireductionism about scientific

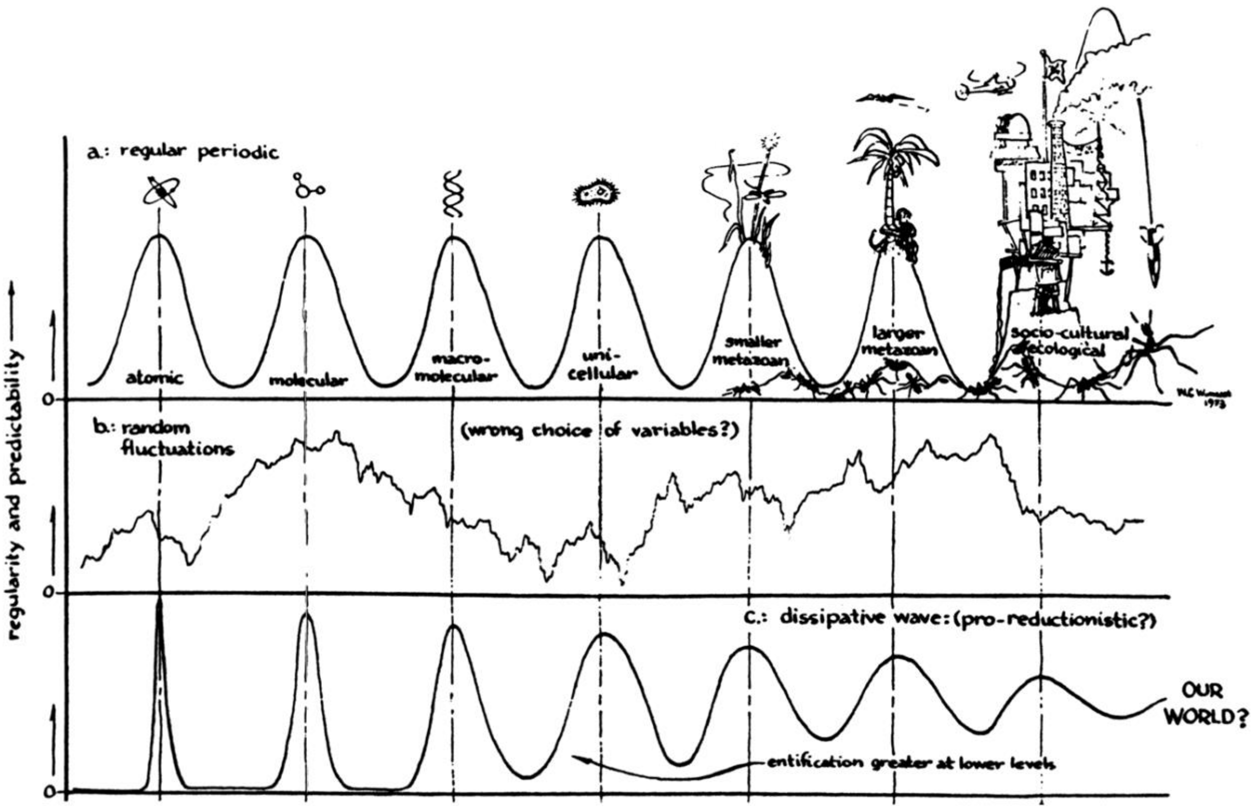

Figure 2: Wimsatt's illustration of levels of organization.

explanation. My survey here has traversed a few different debates and traditions, and it is certainly incomplete. Across the philosophical discussions of levels, several related concepts of levels are variously at play, and often several are invoked without carefully distinguishing among them.[3] There is one important distinction in conceptions of levels I want to point out at this juncture: some are metaphysical, while others are representational. That is, some conceptions of levels regard the entities, properties, or processes in our world—as with claims about levels of organization or high-level causation—whereas others regard our representations of our world—as with claims about the relative abstractness of levels or the levels of our scientific theories. Claims about level relationships among fields of science or of scientific laws seem to lurk somewhere in between metaphysical and

[3] See Potochnik (2017, Chapter 6; 2021) for more on the variety of levels concepts and how these have been conflated.

representational commitments.[4]

Here is an attempt at a generic formulation that perhaps can accommodate all of this variety in the invocation of levels. Most broadly, the antireductionist position about levels of explanation seems to presume that (1) the world (or our representations thereof) is organized into levels, such that (2) our candidate explanations are structured in terms of levels, which motivates (3) the persistence of different fields of science addressing these various levels. This seems to give full voice to the antireductionist impulse, even if my characterization is abstract and rough in order to lump rather than split.

This style of argument for antireductionism has also shown up in philosophical debates about specific scientific investigations, including at least physics, biology, psychology, cognitive science, and the social sciences. The ultimate concerns of these debates vary. In psychology and cognitive science, at issue is whether there is room for mental states as explanations or even as real, causally efficacious components of our world. In physics and biology, the question is how different approaches relate to one another and to other fields entirely, chiefly fundamental physics. In the social sciences, the chief question seems to be social explanations, i.e. explanations citing entities larger than individual agents. Across all these fields, the invocation of levels to counter reductionism also seems to be in service to a stance on methodology or proper modeling approach. Even when what is at stake has varied, these debates have had remarkably similar contours.

[4] A second difference in levels conceptions that can cause confusion is whether (metaphysical) levels are a relation among types or between types and tokens. In this discussion, I presume the former. Even if there is a type-token basis for levels, each token is of many types—i.e. can be categorized with the use of multiple different descriptions—and levels have often been invoked to describe the relationships among those types. Additionally, all fields of science, including fundamental physics, target types rather than tokens.

2. What the Levels Framework Gets Wrong About Explanation

The previous section summarized how a variety of forms of antireductionism about explanation, developed in a variety of philosophical contexts, resist explanatory reductionism by invoking levels. This might seem like an obvious strategy. One might think that all that's needed from the invocation of levels in a project of antireductionism is a way to gesture at the idea that there is something other than the most fundamental: that there is more going on in the world than just microphysical happenings; that our scientific enterprise involves more than fundamental physics; that our scientific laws encompass more than fundamental physical laws; that our scientific explanations come in more varieties than microphysical explanations. But, when one puts this point in terms of levels, much more is taken on board than just this. And, I will suggest in this section, what is taken on board is philosophically problematic. Framing debates about explanation in terms of levels of explanation systematically misconstrues the relationship different explanations bear to one another.

To get at this point, let's start by taking one step back from the levels framework. What is not controversial is that potential explanations come in *varieties*. A variety of physical and chemical theories and models bear on the behavior of gases; genes are investigated with different methods in several subfields of biology; and behavioral phenomena are targeted in studies ranging from neuroscience and molecular genetics to ecological psychology and sociology. Different scientific projects that target the same phenomena generate varieties of potential explanations. I say 'potential explanations' in order to remain neutral on the question of whether all succeed as explanations. Indeed, reductionism is the view that some of these varieties of potential explanations are universally privileged over the others—or will be so at a future stage of science. Antireductionism is the rejection of that universal privilege.

Characterizing varieties of potential explanations as *levels* of explanation entails that the varieties occur in a linear hierarchy or, at the very least, a partial hierarchy. Accounts of levels of explanation tend to presume that potential explanations for some explanandum are either lower level than, higher level than, or at the same level as any other potential explanations (for that explanandum). This just is the idea that explanations are arranged in levels. It is possible, though, that levels of explanation comprise a partial hierarchy: that some potential explanations for a given explanandum are at incomparable levels, while others are related by lower-level-than, higher-level-than, or same-level-as. Invocations of levels of explanation also tend to presume, in line with Fodor's diagram (Figure 1), that there is a many-one relationship among explanations at different levels. Higher-level explanations have often been taken to be more general, i.e., to apply to explananda across a greater range of circumstances, and lower-level explanations to be more specific, i.e. to apply across a more limited range of circumstances (see, e.g., Putnam, 1975; Garfinkel, 1981; Jackson and Pettit, 1992; Sober, 1999). The entities or properties referenced in different levels of explanation are also supposed to bear special relationships to one another: parts and wholes (e.g. Putnam, 1975), sub-components and components of mechanisms (e.g. Craver and Bechtel, 2007), determination (e.g. Yablo, 1992), supervenience (e.g. Kim, 1998), or realization (e.g. Fodor, 1974). Across this variety, it seems higher-level explanations are supposed to reference something that (in one way or another) depends upon what's referenced in lower-level explanations.

But the varieties of potential explanations for a given explanandum are seldom if ever arranged in any of these ways. Consider the three commitments described just above in reverse order. Do varieties of potential explanations reference properties or entities that are connected by metaphysical dependence relations (e.g. composition, determination, supervenience, or realization)? Examples philosophers appeal to in debates about explanatory reductionism tend to have this form. Fodor (1974) contrasts Gresham's Law governing monetary exchanges with imagined lower-level explanations of exchanges of (separately) wampum, dollar bills, and a signed check. Putnam (1975) contrasts the geometric explanation for the square peg not going through round hole with an explanation in terms of the individual atomic structure of the peg and of the edges of the hole. Garfinkel (1981) contrasts the

Lotka-Volterra account of the increasing fox population with an explanation in terms of individual fox and hare births and predation events. An exchange of something valued as currency is required for any monetary exchange; precise atomic structure dictates size relations; individual births and deaths combine to determine population growth. But something is a bit fishy about the lower-level candidate explanations in all of these examples. Namely, these potential explanations are not actually types of explanation generated in scientific research. Where is the scientific research on wampum exchanges (or signed-check exchanges), on the precise atomic structure of a specific one-inch sided peg, or on the births and deaths of foxes and hares in a specific population? Without the existence of such research, we can't expect these examples to inform our judgments on how to choose among a variety of potential explanations.

In contrast, in situ examples of potential explanations for a given explanandum do not tend to have this form. Behavioral phenomena—say, a tendency to heightened aggression—are investigated in a range of different fields (Longino, 2013). Neuroscience identifies neural structures and pathways associated with this tendency; molecular genetics identifies genes associated with the tendency; ecological psychology and sociology identify social and environmental influences. None of these explanations for heightened aggression is related to another by composition, determination, supervenience, or realization. Rather, these different explanatory strategies specify different causal influences on the behavioral phenomenon in question (which may or may not also bear causal relationships to one another). Genes can causally influence neural structures and pathways but do not compose or realize them, while social and environmental influences are separate influences that can also have neural effects. This result is easy to replicate with other instances of multiple investigations targeting the same phenomenon. Potochnik (2010) argues this is the general form of the relationship between competing "levels" of explanation in science, as different potential explanations broadly fail to be related by metaphysical determination.[5] Ylikoski (2014) argues that, in the social sciences, there is a variety of micro- and macro- social explanations that are not arranged in levels.

At first glance, it seems genetics might fare better, with the anticipated dependence relations connecting explanations in molecular genetics to those in classical genetics. These investigations invoke different specifications (structural and functional, roughly) of the very same entities (genes). But the styles of explanations formulated in molecular genetics and classical genetics do not capitalize on that relationship. Rather, molecular genetics provides information about molecular genes associated with some trait of interest, often via genome-wide association studies. Very seldom is the causal role of any given molecular gene able to be identified. And classical genetic explanations, as in behavioral genetics, partition overall influence on a trait into genetic heritability vs. environmental influences (Longino, 2013). So, despite the apparent promise, molecular genetics and classical genetics are not well positioned to provide candidate explanations citing entities or properties related by metaphysical dependence. The upshot of this discussion is that entities and properties featured in a variety of potential scientific explanations for the same or related phenomena in scientific research generally bear no special metaphysical dependence relationship (such as composition, realization, determination, or supervenience) to one another.

This finding also interferes with the expectation that higher-level explanations are more general than lower-level explanations, as that expectation of relative generality issues from the anticipation of a many-one relationship between the relevant entities or properties stemming frommultiple realization, supervenience, determination, or composition. There is also a deeper problem beyond this. The generality of an explanation depends on its degree of abstractness, i.e. how many details are specified. Specifying fewer details results in a more general account, while incorporating more details decreases

[5] Franklin-Hall (2016) calls this relationship 'horizontal' rather than 'vertical'; terminology that reinforces an expectation of dependence of higher-level wholes on their lower-level parts.

the generality (and increases the precision) of an account. But degree of abstractness is a property of representations, or characterizations, not of what is being represented or characterized. Generality, as in scope of applicability, might be a metaphysical property, but the relative generality of an explanation is influenced by representational choices, namely what to include or exclude from the explanation.

Two implications of this are important for present purposes. First, the abstractness or generality of an explanation needn't relate to the metaphysical dependence relations typically thought to characterize levels, such as realization, supervenience, or composition (Potochnik, 2021).[6] By specifying additional properties or omitting mention of other factors (e.g.), one can make an explanation more or less general regardless of what entities and properties it features. Some explanations in microphysics are highly general, while others apply only in a finely specified set of conditions. The same goes for biology and economics. Discussions of levels of explanation have generally presumed that lower-level is more specific and higher-level more genera, but in other contexts philosophers have regularly touted the relative generality of accounts in physics vs. the so-called special sciences. Thus, the explanatory value of generality is not necessarily a point in favor of antireduction about explanation, and specifying additional explanatory detail need not result in a reductive explanation.

The second implication of abstractness being representational that I want to emphasize is that it is quite common to have *incommensurate* degrees of generality—i.e. two representations that cannot be ranked in their generality but simply are general in different respects. Abstractness (and thus generality) is achieved by omitting details. Omit or include details about different things, and the resulting representations are of incommensurate generality. They specify different aspects of the world in virtue of what is depicted and generalize to different ranges of circumstances in virtue of what is

[6] The determination relation seems more closely related to generality but also more distantly related to levels of explanation as they have often been understood. Franklin-Hall (2016) acknowledges this by pointing out the divergence between determinate/determinable and micro/macro yet persists in calling the determinate/determinable relation "vertical"—which to my mind continues the conflation of fineness of specification and compositional determination that I aim to disambiguate here.

omitted. Philosophers have debated the proper degree and variety of generality in our scientific explanations, but that is incidental to the present point. The point here is simply that varieties of potential explanations quite often cannot be ranked by degree of generality, so this feature of the levels of explanation framing fails to obtain on a regular basis.

Finally, let's consider the very basis of framing varieties of potential explanations as *levels* of explanation: whether potential explanations for a given explanandum are arranged in a linear hierarchy or partial hierarchy. The delineation of levels requires that potential explanations be sortable into lower-level-than, higher-level-than, or same-level-as; if one anticipates a partial ordering rather than linear hierarchy, a fourth category of "incommensurate level" is also available. The argument above that potential scientific explanations of the same phenomenon often are not related by citing entities or properties related straightforwardly by composition, realization, supervenience, or determination already suggests difficulties with sorting potential explanations into a linear hierarchy, as those are typically cited as the basis for delineating levels. Even if one aims for a partial ordering, the enormity of the "incommensurate level" category is troubling, if—as I have argued—most or all potential explanations for the same explanandum are not related by composition, realization, supervenience, or determination. Appeal to levels of explanation was meant to categorize potential explanations in an informative way, but, for many of our candidate scientific explanations, even with the same explanandum, the anticipated means for sorting into levels are unavailable.

The problem goes deeper. I have skirted this issue so far in this chapter, but delineation of levels of explanation has typically relied on delineating levels of organization based on relations like material and mechanistic composition, spatial and temporal scale, and realization But, as I have explored elsewhere, it turns out that these relations together do not determine a linear hierarchy or useful partial ordering (Potochnik, 2017, Chapter 6) and any individual relation among them cannot be used separately to determine a linear hierarchy or useful partial ordering (Potochnik, 2021). I want to emphasize that the problem is not that there are occasional exceptions to an ordering or orderings that cannot be universal in scope. Rather, it seems our world—or at least the properties and entities into which science has carved it—simply is not composed of levels (see also Thomasson, 2014; Eronen, 2015). Indeed, when Kim (2002)—a philosopher perhaps best known for a causal exclusion argument that presumes the levels framework—explored the basis for levels, even he came up short.

The commitments I take the levels framework to have and the difficulties I have pointed out with each of these commitments are summarized in Table 1. If levels of organization cannot be used to impose a linear hierarchy on our explanations, and generality rankings of our varieties of explanation do not result in a linear hierarchy, and the varieties of potential scientific explanations we observe in situ do not seem to bear any of the anticipated hierarchical relationships to one another, then I am not sure what the basis would be for the presumption that scientific explanations come in levels.

Commitments of the view that explanations come in levels	Difficulties with these commitments
Explanation varieties arranged in linear hierarchy	Incommensurate rankings are commonplace
Higher-level explanations are more general than lower-level explanations	Abstractness is a representational, not metaphysical, property
Higher-level explanations cite something that depends on what lower-level explanations cite	Potential explanations often bear no special relationship to one another

Table 1: Commitments of the view that potential explanations come in levels and the difficulties facing them

It follows from all of this that we should pause to consider the proper framing before employing the seemingly straightforward diagrams of level relations that recurs in many discussions of levels of explanation and causation. Figure 1 depicts an important example of such a diagram; another primary example is the diagram commonly used to depict the causal exclusion argument that originated with Kim (e.g. 1998). The vertical lines illustrating realization or supervenience may seem to be uncontroversial in light of the broad acceptance of physicalism and material composition, but these commitments do not suffice as grounds for the assertion that our scientific explanations or the causal relationships they feature bear the implied metaphysical relationships to one another. Those vertical lines must be earned rather than assumed, lest the very framing of the question inherit our mistaken assumptions about the variety of potential scientific explanations.

Philosophical discussions about explanatory reductionism have, by and large, presumed explanations come in levels, identified candidate explanations on the basis of that expectation, and then assessed quality and status of those candidate explanations to make a determination regarding reductionism. The debate is transformed if instead we look to science to see what varieties of potential explanations for the same phenomena are identified and what relationship those bear to one another. One favoring the standard strategy might argue that that approach gets at a metaphysically deeper picture of alternative explanations or that, in the fullness of time, varieties of scientific explanation will tend toward the predicted relationship of levels. In response to the former, I'll point out that so long as the target of our philosophical accounts of explanation is *scientific* explanations, the better strategy is one that applies to explanations actually formulated in science. In response to the latter, I see no reason to expect that science is moving toward a division of labor ordered by metaphysical relationships like composition, realization, or supervenience. This is supported by considerations like those I have already raised in this section with framing antireductionism about explanation in terms of levels.

3. Antireductionism Without Levels

Conceptualizing antireductionism in terms of levels systematically misconstrues the relationship different candidate explanations bear to one another. Candidate explanations do seem to relate many-one to their explananda, as anticipated with the levels framing. Potential explanations come in varieties. But, as discussed in the previous section, "levels" is not an apt description for those varieties. The relationship potential explanations of the same phenomena bear to one another is not a linear or partial hierarchy orderable by generality and defined by metaphysical dependence of the featured entities and properties. In this section, I use the shortcomings of the levels framing to inspire an alternative approach to antireductionism. This alternative approach better describes the variety of candidate explanations we see in scientific research and the relationships these explanations bear to one another, and it also better accounts for the considerations that guide the selection among candidate explanations. I discuss these three significances below. I conclude this section by pointing to some problematic downstream implications to which framing antireductionism in terms of levels has given rise and, accordingly, an antireductionism without levels helps us avoid.

First, having jettisoned the expectations that accompanied the levels framing, let's reconsider how different potential explanations of the same explananda relate to one another. In Section 2, I anticipated an alternative to straightforward metaphysical dependence (whether composition, determination, supervenience, or realization): that different explanatory strategies specify different causal influences on the phenomenon in question, influences that may or may not also bear causal relationships to one another (see also Potochnik, 2010).[7] Just as attention to in situ varieties of scientific explanations revealed that the levels framing is frequently inapt, this also lends prima facie support to this alternative framing. Consider again the variety of investigations that aim to explain human behavioral tendencies. These can feature (at least) molecular genes, neurological features, environmental influences, social context, and more. These factors interact in their influence on human behavior, and some also causally influence one another, as with molecular genes' and the environment's impact on neurological development.

This example supports an additional consideration in favor of the expectation of different potential explanations targeting distinct influences on a phenomenon. In a wide variety of scientific research, factors that are non-local turn out to be key influences on phenomena. This amounts to an empirical vindication of at least one form of antireductionism: it turns out that large-scale influences,

[7] I suspect most if not all scientific explanations include causal information, but what I say here is not intended to commit one to the view that all scientific explanations are causal explanations. There might be exceptions to

distal influences, and structural influences regularly shape the happenings in our world. Examples are easy to generate. In ecology, abundance (i.e. population sizes) traditionally was thought to be determined locally by interactions with competitors but is now recognized to be shaped globally, such as in the evolution of specialists and generalists (e.g. Gaston and Blackburn, 2000). Dynamical systems theory has been fruitfully applied to research ranging from physics to ecology, cancer, and cognitive science. And it is now widely appreciated that mitigating racism involves not just changing minds but renovating social systems. The potential for significant non-local influence means there are more places to look for explanations and less reason to think independently generated explanations bear any metaphysically deep relationships to one another, such as determinables and determinates or realized and realizers. Phenomena in our world are shaped by so many different influences, operating at different timescales and spatial scales, that the potential explanations for one of those phenomena seldom are different characterizations of the very same states of affairs. Individuals' racist views may causally influence the features of social systems, and similarly for interspecific competition and evolutionary trajectories, but in neither case do the former compose, realize, or determine the latter.

Debates about explanatory reduction have tended to conflate two questions, roughly: (1) how finely (i.e. at what "level") explanatory factors should be characterized, and (2) whether explanatory factors tend to be local, perhaps even components of the system exhibiting the phenomenon to be explained (i.e. components at "lower levels"). Scientific research has empirically shown (2) to be wrong: explanations cite the distant, the largescale, and the structural in order to shed light on a variety of phenomena. This outcome should lead us to shift our gaze outward rather than down, so to speak, when looking for candidate explanations. And that, in turn, leads the first question of reductionism, (1) above, to seem rather beside the point. If entirely different factors are targeted in different candidate explanations, then the question of how finely to characterize a factor does not arise. This relates to the point I made in the previous section that, when we consider candidate explanations actually generated in scientific research, they turn out not to be clearly related by determination, realization, supervenience, or composition.

Second, the framing of levels also suggests candidate explanations are less numerous than they in fact are. Looking merely to components on lower organizational levels suggests we will have, at most, the number of explanations as there are levels; for instance, the social explanation, psychological explanation, neurological explanation, and genetic explanation. When we shift our gaze outward instead of down, as I suggested just above, shifting our antireductionist expectation from levels to varieties of influences, this opens the door to the recognition of a much wider variety of candidate explanations. Candidate explanations, it seems, may be as numerous as factors that significantly bear on the phenomenon—or even as numerous as various partially overlapping sets of these factors that may be targeted in different investigations. Such candidate explanations differ not just in what factors are cited but also (as a result) in what circumstances or to what varieties of phenomena they apply. This is anticipated by what I said in Section 2 about representations with incommensurate generality, that is, that generalize across different ranges of circumstances. Indeed, this is hardly surprising when we take into account the different research projects within which different explanations are formulated. Positing that there are potential explanations at different levels does nothing to resolve which of *these* potential explanations are better (and in which circumstances). And, then, the need for resolution on this broader question renders the question of better level of explanation rather redundant.

For example, explaining some phenotypic trait, say, variation in coloration in Harris sparrows, can take place in the context of research into frequency-dependent selection, explaining this as an instance of the hawk-dove game dynamics (Maynard Smith, 1984). Or an explanation may be generated in research on phenotypic plasticity, with an explanation that bears on this as an instance of environmental influence on trait development. (See Potochnik, 2016, for a more extended discussion of this example.) There are many more possibilities beyond these two: explaining trait variation within a population is of interest in a number of biology research programs. Oftentimes the difference is associated with differently characterized explananda, but specification of explanandum isn't sufficient to single out just one—or even a few—explanations (Potochnik, 2016). And, nothing is special about this example. Relevance across multiple research projects and variable significance for those research projects is common for phenomena scientists aim to explain.

Third, this wide variety of candidate explanations and how they relate also complicates the grounds for deciding among the candidates. Classically, for antireductionism couched in terms of levels of explanation, advantages like generality, breadth, and stability have been touted as grounds for preferring non-reductive explanations. (For some recent discussions, see Weslake, 2010; Blanchard et al., 2018; Bradley, 2020.) I have suggested that, if varieties of explanation do not come in levels and do not bear special relationships to one another, then candidate explanations cannot be straightforwardly ordered with regards to generality. For antireductionism without levels, different varieties of generality suit explanations to contribute to different research projects, with different aims. We should not expect an across-the-board ordering for any other measures that may be relevant to an explanation's quality. Similar considerations may still play a role in determining which candidate explanation(s) fits the bill. But, given the ease of generating potential explanations with different forms of generality (i.e. that generalize to different ranges of systems), it is possible or even likely that such considerations will vary with the requirements of different research projects. The question may not be which explanation is most general, stable, or offers greatest breadth or guidance but, rather, which explanation has these properties in the right combination and regarding the right features to be most valuable to a specific research project.

If this is so, then this suggests a form of explanatory pluralism. One single explanation may not win out against all other candidates, but rather multiple explanations may be developed across science, each of which best addresses some research needs but not others. This is different from an explanatory pluralism developed within the levels framework, such as by Jackson and Pettit (1992) and Sober (1999), as those views adopt the expectation of levels of explanation ordered by relative generality and embrace pluralism with regards to how much generality is desirable. Note that one might follow my urging to reframe antireductionism without a commitment to such an explanatory pluralism: one may hold that there is always a single, best explanation that is non-reductive in the ways I have outlined (perhaps a single, integrative explanation that draws from all relevant research projects). On the other hand, if one does accept explanatory pluralism of this form, then this opens up a significant role for scientists' interests and priorities in shaping the nature of scientific explanations—due to what they emphasize and what they sideline in their particular research projects. This is explored for varieties of explanation in cognitive science by Potochnik and Sanches de Oliveira (2020), who call this different "explanatory styles".

To summarize, a better antireductionism about explanation stems from the insights that different potential explanations regularly feature entirely different factors influencing the phenomenon, that these potential explanations vary in what they attend to and what they abstract from as they are developed in and contribute to different research programs, and that grounds for deciding among these potential explanations include considerations that may also vary with different research programs and perhaps even with something as basic as scientists' interests. The antireductionism comes in granting the legitimacy or even preferability of at least some of these potential explanations that do not feature local microphysical happenings. Explanatory pluralism results from additionally asserting that multiple of these potential explanations are warranted for a single phenomenon (as characterized in some explanandum).

Nowhere in these statements of explanatory antireductionism and explanatory pluralism is reference to levels needed, and, I propose, such reference would actually be a liability. An antireductionism based on levels fails to incorporate these features and is impoverished by their absence. Further, even if the points made in this section are somehow accommodated, the levels framing remains problematic for philosophical debates about scientific explanation. Such a framing easily can slide into the presumptions I argued against in Section 2. This framing is also associated with other problematic and unearned ideas. For one, antireductionism based on levels of explanation has been taken to suggest that explanandum and explanans should be on the same "level," that is, regard similarly sized objects operating at similar timescales (e.g. Wimsatt, 1972; 2007). This may work as a defense against reductionism, but it also defines away the possibility of large-scale and structural causes—and, for that matter, the possibility of tiny entities sometimes wielding great explanatory power. An instance of this is individualism in social science, where behavior is expected to be fully explained by the properties of individuals; see Haslanger (2015) for an argument against individualism in favor of structural explanation, or explaining behavior with reference to systems in which individuals participate. The expectation that explanations should match the level of what is being explained is clearly wrong. As discussed above, scientific explanations regularly cite the distant, the largescale, and the structural in order to shed light on a variety of phenomena, as with structural explanation in the social sciences. It is a further liability of the levels framework that it obscures this in order to counter the view that all explanations trace back to microphysical happenings.

4. A New Working Hypothesis about Scientific Explanation

Oppenheim and Putnam's (1954) stance and Fodor's (1974) rebuttal were both explicitly formulated as 'working hypotheses' about how the relevance of levels will play out in science: the former reductive unity, the latter independence of levels of realization. Both working hypotheses, I submit, have been proven wrong by scientific advances. At this point, there is ample scientific evidence in favor of antireductionism about scientific explanations. With few if any exceptions, the so-called special sciences continue about their business, indifferent to any breakthroughs in microphysics, and the explanations they produce are not treated as provisional, awaiting vindication by reduction. But, just as importantly, there is also ample scientific evidence that explanations don't come in levels. Different fields and subfields that target the same phenomena focus on different factors that by and large bear no straightforward metaphysical relationship to one another, and largescale and systemic factors can be key to explaining many phenomena.

Both previous working hypotheses—reductionism and levels of explanation—share starting presumptions about how our world operates that turn out to be wrong. Both of these philosophical positions presume that the key to explaining phenomena is their features, their immediate causes, and perhaps what composes them; we might call this 'localism' about explanation. Thus, the choice in explanation has been framed as between on-site microphysical happenings or lumpier characterizations of those happenings. But localism about explanation is wrong. Phenomena are regularly determined by largescale and distant factors, by structural and contextual factors, by systems in which they participate. Abundant scientific and philosophical research supports this claim, including complexity research such as the aforementioned dynamical systems theory, developmental systems theory, systems biology, and network theory. Recall from above that one of the entry points for the levels framework in philosophy was in Herbert Simon's work on complexity. Historically, at least in that tradition, positing non-reductive levels of explanation was a way to accommodate complexity. But since then, the antireductionist levels framework and complexity research have parted ways: the former has retained a commitment to localism about explanation, while the latter is predicated on its rejection. And rightly so, it seems to me. Localism about explanation is demonstrably false.

This inspires the new working hypothesis I propose about scientific explanation. In accordance with the view outlined in Section 3, I propose that prospective explanations are many and varied, often including some featuring large-scale, distant, and structural influences. The decision is not about how fine-grained our characterization of local factors should be but rather about which factors at what scales we should attend to. This may well have at least some objective determiners, but I suspect some of the determination will be left to what scientists and their audiences prioritize, intentionally or not, via the specific research projects scientists pursue. This characterization fits better than either reductionism or antireductionist levels with the variety of potential explanations encountered in scientific research and with how those candidate explanations relate to one another.

The two diagrams of levels featured in above figures have held remarkable sway over our field, so I have tried to offer a competing image in Figure 3. It is more mundane: a failure to draw lines to demarcate levels or arrows to demarcate metaphysical determination relationships just shows up as blank page. Aside from what this image does not include, the important features are (1) that the explanandum is grouped with multiple different sets of related phenomena, (2) that those groupings are associated with the identification of different explanatory factors, and (3) these different explanatory factors bear no special relationship to one another (at least not in general). E is the explanandum; X, Y, and Z are factors that comprise candidate explanations for E.

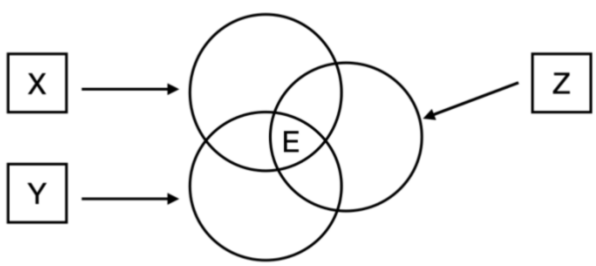

Figure 3: Antireductionism of explanation as a thesis about different explanatory factors operating at different scales

Here is a brief example to illustrate this characterization. What explains the scarlet ibis's bright coloration? (We could additionally specify the contrast class: rather than the white feathers of the closely related white ibis.) One candidate explanation focuses on the scarlet ibis's ability to metabolize carotenoids. This highlights a primary form of coloration across bird species. Another candidate explanation focuses on the specific carotenoid carrier protein found in the scarlet ibis's blood. Yet another candidate explanation postulates the role of the scarlet ibis's vibrant coloration in mate attraction. Each of these is the subject of scientific research. Each distinguishes the scarlet ibis's coloration from that of the white ibis. Each casts light on a different range of related phenomena: from avian coloration in general, to the scarlet ibis's particular metabolism, to the role of bright coloration in sexual selection. See Figure 4. They are not competing explanations. One or another may turn out not to be exactly right, but it's possible that all are correct. Depending on the specifics of our account of explanation, we may require more to be said about one or another for it to count as an explanation, favor one over others, or even anticipate their integration in a single explanation. However, I think the most likely outcome—and what best accommodates the realities of scientific investigation—is that all three of these explanations are accepted (or suitably refined versions if new evidence comes to light). But regardless, the point for present purposes is that the choice among potential explanations is not how finely to characterize the local details but rather which kinds of factors the investigation should target. Reductionism about explanation is incorrect—and so is antireductionism that relies upon levels, and the localism it presumes.

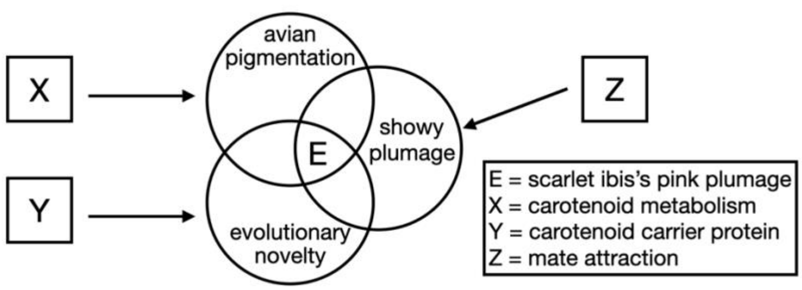

Figure 4: Coloration of scarlet ibis illustration of different explanatory factors operating at different scales

Acknowledgments

Thanks to Alistair Wilson and Katie Robertson for inviting my participation in this volume, as well as to Sahar Heydari Fard, Sally Haslanger, and the Structures Reading Group for helpful comments on an earlier draft.

References

Blanchard, T., Vasilyeva, N. and Lombrozo, T. (2018) “Stability, Breadth and Guidance.” *Philosophical Studies* 175(9): 2263-2283.

Bradley, D. (2020) “Should Explanations Omit the Details?,” *The British Journal for Philosophy of Science* 71 (3):827-853.

Clarke, C. (2016) “The explanatory virtue of abstracting away from idiosyncratic and messy detail.” *Philosophical Studies* 173(6): 1429-1449.

Craver, 2007. *Explaining the Brain: Mechanisms and the Mosaic Unity of Neuroscience*. London, UK: Oxford University Press.

Craver, C.F. and Bechtel, W. 2007. “Top-Down Causation Without Top-Down Causes.” *Biology & Philosophy* 22: 547-563.

Eronen, M. 2015. “Levels of organization: a deflationary account.” *Biology & Philosophy* 30: 39-58.

Fodor, J.A. 1974. “Special Sciences; or, The Disunity of Science as a Working Hypothesis.” *Synthese* 28(2): 97–115.

Franklin-Hall, L.R. 2016. “High-Level Explanation and the Interventionist’s ‘Variables Problem’.” *The British Journal for the Philosophy of Science* 67(2): 553–577.

Garfinkel, A. 1981. *Forms of Explanation: Rethinking the Questions in Social Theory*. New Haven, CT: Yale University Press.

Gaston, Kevin J., and Tim M. Blackburn. 2000. *Pattern and Process in Macroecology*. Oxford: Blackwell

Scientific.

Haslanger, S. 2016. “What is a (Social) Structural Explanation?” *Philosophical Studies* 173: 113–130.

Haug, M. C. 2011. “Abstraction and Explanatory Relevance; or, Why Do the Special Sciences Exist?” *Philosophy of Science* 78(5): 1143–1155.

Hempel, C. G., and P. Oppenheim. 1948. “Studies in the Logic of Explanation.” *Philosophy of Science* 15(2): 135–75

Jackson, F., and P. Pettit. 1992. “In Defense of Explanatory Ecumenism.” *Economics and Philosophy* 8(1): 1–21.

Kim, J. 1998. *Mind in a Physical World: An Essay on the Mind-Body Problem and Mental Causation*. MIT Press.

Kim, J. 2002. “The layered model: Metaphysical considerations.” *Philosophical Explorations* 5: 2–20.

Longino, H. 2013. *Studying Human Behavior: How Scientists Investigate Aggression and Sexuality*. Chicago: University of Chicago Press.

Maynard Smith, J. 1982. *Evolution and the Theory of Games.* Cambridge UP.

Oppenheim, P. and H. Putnam, 1958, “The unity of science as a working hypothesis”, in H. Feigl et al. (eds.), Minnesota Studies in the Philosophy of Science, vol. 2, Minneapolis: Minnesota University Press.

Potochnik, 2010. “Levels of Explanation Reconceived.” *Philosophy of Science* 77: 59-72.

Potochnik, 2011. “A Neurathian Conception of the Unity of Science.” *Erkenntnis* 34: 305-319.

Potochnik, 2016. “Scientific Explanation: Putting Communication First.” *Philosophy of Science* 83: 721—732.

Potochnik, 2017. *Idealization and the Aims of Science*. Chicago: University of Chicago Press.

Potochnik, A. and Sanches de Oliveira, G. 2020. “Patterns in Cognitive Phenomena and Pluralism of Explanatory Styles.” *Topics in Cognitive Science* 12: 1306-1320.

Potochnik, A. 2021. “Our World Isn’t Organized into Levels,” in D. S. Brooks, J. DiFrisco, and W. C. Wimsatt (Eds.), *Levels of Organization in the Biological Sciences*, MIT Press.

Putnam, H. 1975. *Philosophy and our Mental Life*. Cambridge: Cambridge University Press.

Simon, H. 1962. “The Architecture of Complexity.” *Proceedings of the American Philosophical Society* 106: 467-482.

Sober, E. 1999. “The Multiple Realizability Argument against Reductionism.” *Philosophy of Science* 66(4): 542–64.

Thomasson, A. 2014. “It’s a Jumble Out There.” *American Philosophical Quarterly* 51: 285-296.

Weslake, B. 2010. “Explanatory depth.” *Philosophy of Science* 77(2), 273-294. Chicago.

Wilson, J. (2013). “Metaphysical Emergence: Weak and Strong.” In Stephen Mumford & Matthew Tugby (eds.), *Metaphysics and Science*. Oxford UP.

Wimsatt, W. 1972. “Complexity and Organization.” *PSA: Proceedings of the Biennial meeting of the Philosophy of Science Association*, K. F. Schaffner and R. S. Cohen (Eds.), D. Reidel.

Wimsatt, W. 2007. *Re-Engineering Philosophy for Limited Beings*. Cambridge MA: Harvard University Press.

Yablo, S. 1992. “Mental Causation.” *The Philosophical Review* 101(2): 245–80.

Ylikowski, P. 2014. “Rethinking Micro-Macro Relations.” *Rethinking the Individualism-Holism Debate: Essays in the Philosophy of Social Science*, J. Zahle and F. Collin (Eds.), Springer.