CHAPTER 22

LEVELS, INDIVIDUAL VARIATION, AND MASSIVE MULTIPLE REALIZATION IN NEUROBIOLOGY

KENNETH AIZAWA AND CARL GILLETT

> No one supposes that all the individuals of the same species are cast in the very same mould. These individual differences are highly important for us, as they afford materials for natural selection to accumulate.... These individual differences generally affect what naturalists consider unimportant parts; but I could show by a long catalogue of facts, that parts which must be called important, whether viewed under a physiological or classificatory point of view, sometimes vary in the individuals of the same species.
>
> Charles Darwin, *On the Origin of Species*

uroscientists, like all biologists, hold two fundamental beliefs about nervous systems. st, they believe that nervous systems can be studied at any number of distinct but erdependent levels of organization in which entities at one level are explained by qualitatively different entities at one or more lower levels that are taken to compose m. Neuroscientists study structures as large as communities of interacting organisms d as small as individual proteins. There are thus a number of neurobiological levels.

The second fundamental belief is that nervous systems display individual variation. Subsequent research has shown that Darwin (1964) surely understated the case, especially with the subject matter of the neurosciences, when he observed that not all individuals of a species are cast from the same mold. Organisms obviously vary in their genetic makeup, but given distinct histories of interaction with their environments even genetically identical individuals will diverge in their phenotypic details. In truth, no two organisms are exactly alike, molecule for molecule, cell for cell, or organ for organ—especially when the molecules, cells, and organs in question are those studied by the neurosciences.

Combining these two fundamental beliefs, we may say that as far as can currently be determined, individual variation appears at every level of neurobiological organization. As a result, because component entities such as realizer properties vary at particular levels, we contend that we have overwhelming scientific evidence for what we call the massive multiple realization (MMR) hypothesis about psychological properties:

> (MMR) Many human psychological properties are multiply realized at many neurobiological levels.

Putting the thesis in other words, MMR is the claim that for many human psychological properties, the instances of these properties are realized by different lower level properties at many of the levels studied in neuroscience.[1]

As our opening points suggest, and as the evidence we highlight supports, the MMR hypothesis is uncontentious for many working neuroscientists (although they obviously do not refer to the relevant phenomena in the terms used in the thesis). In contrast, the existence of *any* multiple realization of psychological properties by neuroscientific properties, let alone *pervasive* or *massive* multiple realization, has been bitterly fought over by philosophers. To understand these differing reactions of neuroscientists and philosophers, it is important to briefly lay out the recent background to debates in general philosophy of science, the philosophy of psychology, and the philosophy of neuroscience. Setting the scene in this manner allows us to better situate our work in the chapter and the overall position we ultimately defend in relation to recent philosophical battles.

All areas have narratives (stories, if you prefer) about the present issues and competing positions. Though obviously a caricature, the following is hopefully a useful sketch of one common narrative current in much philosophy of neuroscience about the recent dialectical state-of-play and philosophical battle lines. (We should note that researchers in philosophy of psychology obviously have different stories to tell, but one of these narratives does reflect this understanding of the debate.[2])

On one side, so the story goes, we find proponents of cognitive science (the name of Jerry Fodor is often dropped at this point) who are taken to endorse the existence of multiple realization. These defenders of cognitive science are also taken to use multiple realization to establish the autonomy of cognitive science from neuroscience, where the latter is read as the claim that neuroscience and cognitive science do not intertheoretically constrain each other. On the other side, the story continues, we find those who emphasize the importance of neuroscience and defend the existence of intertheoretic constraints between cognitive sciences and neuroscience, and who consequently use such intertheoretic constraint to attack the existence of multiple realization. (Writers offering such arguments include Bechtel and Mundale, 1999, Shapiro, 2004, and others.) Along with these opposing commitments, our two camps are also read as having conflicting views about the possibility of Nagelian reduction (Nagel, 1961), and the existence of univocal realizations and/or species-specific identities between neuroscientific and psychological properties. The defenders of cognitive science are taken to reject such claims, and those sympathetic to neuroscience are interpreted as defending them.

This narrative obviously posits a range of ongoing disputes that have implications far beyond the philosophy of psychology and the philosophy of neuroscience, because the questions putatively at issue concern the status and importance of various scientific disciplines. It is thus only a small step from these scholarly discussions about the nature and appropriate relations of neuroscience and psychology, and their respective entities, to more pragmatic debates over the appropriate funding levels for these scientific areas and particular approaches within them. Unsurprisingly, as is often the case when funding discussions become public, the resulting debates in philosophy have been heated and hard-fought. Our goal in this chapter is to engage these philosophical disputes over multiple realization from some fresh directions and attempt to reconnect the concerns of philosophers with the frameworks that working neuroscientists take to be mundane.

First, though passion has not been lacking in recent discussions in philosophy, what has been missing is any precise philosophical framework for a key element of these debates in the compositional relations between the levels of entities in neuroscience, including realization relations between properties. Our initial attempt to freshen the recent debates focuses on addressing this deficit. We begin by using a concrete, well-understood case from neurobiology, in section 1, to highlight variation and levels in neurobiology and also to sketch the general nature of the concepts of composition routinely posited in explanations in the sciences. Using these more general observations as a platform, we then provide precise theory schemata for both the realization relations between properties and multiple realization itself.[3]

Our framework for realization and multiple realization provides new theoretical resources, and we also seek to freshen the debate in a second way by using our framework to examine a selection of empirical evidence to highlight the nature of a number of neurobiological levels. We therefore give a brief sampling of scientific findings in section 2, illuminating the variety of such levels, and show that there is plausibly important individual variation at every physiologically significant level of organization in the nervous system—from proteins to whole brains. Applying this theoretical work on realization and multiple realization, we consequently show that such evidence about individual variation provides a prima facie plausible case for MMR. Our more detailed theoretical frameworks for scientific composition thus illuminate why working scientists apparently find multiple realization, though described in different terms, to be so mundane.

Since so many philosophers have thought that multiple realization is far from trivial, perhaps even being scientifically damaging, we finish, in section 3, by exploring philosophical concerns about the MMR hypothesis. We show that our more precise theoretical framework for realization, in combination with neurobiological evidence, establishes that a range of common objections to the existence of multiple realization are mistaken. For example, we show that multiple realization simply does not establish the methodological autonomy of cognitive science and other branches of psychology, but actually supports the utility of a coevolutionary research strategy based around methodological interactions between the psychological and neurobiological sciences.

One of our goals in the chapter is therefore to show that the lack of a theoretical framework for scientific composition has been highly damaging, because we demonstrate that with a precise account of realization relations in the sciences one can establish the error of *both* of the sides commonly taken to be battling in recent philosophical debates. With better accounts of scientific composition, realization, and multiple realization in hand, we show that the empirical evidence underpinning the standard neuroscientific belief that nervous systems have individual variation at many levels of organization supports *both* MMR *and* intertheoretic constraint between cognitive science and neuroscience. As we suggested, such a combination of multiple realization and methodological interaction between neuroscience and psychology has been anathema to many philosophers, though it appears mundane to working scientists in both disciplines. Our hope is that getting clearer about scientific composition generally, and realization and multiple realization in particular, restores a balance between the outlooks of philosophers and neuroscientists, not least by challenging a number of mistaken and damaging positions that have recently taken root in the philosophies of neuroscience and psychology.

1. COMPOSITION IN THE SCIENCES: UNDERSTANDING REALIZATION AND MULTIPLE REALIZATION

In this section we seek to give a clearer picture of the compositional concepts posited in mechanistic explanations in the sciences, some of which are summarized in table 22.1.[4]

We should remark that terms like *realization*, *constitution*, and *implementation* have been used in all manner of ways by theoreticians, whether metaphysicians, logicians, or philosophers of science. For example, the word *realization* has been used by philosophers and scientists to refer to a number of very different concepts

Table 22.1. Compositional Relations in the Sciences

Type of Entity as Relata	Compositional Relation
Processes	Lower level processes *implement* a higher level process
Individuals	Lower level individuals *constitute* a higher level individual
Properties	Lower level properties *realize* a higher level property
Powers	Lower level powers *comprise* a higher level power

in a range of distinct projects. However, given the focus here, we exclusively use these terms to refer to the relevant compositional relations posited in the sciences and thus offer a view of realization that seeks to capture the compositional relations between properties posited in mechanistic explanations.[5]

Although we are concerned with scientific compositional relations in general, we focus most of our attention on individuals and properties and their compositional relations in constitution and realization. In treating properties, we assume a weak version of the "causal theory of properties." This is a variant of Shoemaker's (1980) account under which a property is individuated by the causal powers it *potentially* contributes to the individuals in which it is instantiated. On this view, two properties are different when they contribute different powers under the same conditions.

To concretely anchor our work and aid the explication of our accounts of realization and multiple realization, we focus on a familiar case from neurobiology, where our explanations are well confirmed, in recent mechanistic explanations of color processing in the human retina at a number of neurobiological levels. The sciences provide mechanistic explanations of the retina that take it to be constituted by individuals at cellular, biochemical, and atomic levels and take the chromatic processing properties of the human retina to be correspondingly realized by properties and relations at the cellular, biochemical, and atomic levels, among others.

Focusing on individuals (as shown in figure 22.1), the sciences now take the retina to be constituted by, among other things, rods and cones; take rods and cones to be constituted by, among other things, complex light-sensitive protein molecules; and take such molecules of photopigment to be constituted by various atoms. Turning to properties and relations, as we relate in more detail as we progress through this section, the sciences also provide mechanistic explanations of the properties of the individuals at higher levels in terms of the properties of individuals at lower levels. For instance, the sciences take the retina's property of processing color to be realized by, among other properties/relations, the light absorbing and signaling properties of retinal cells and their pattern of synaptic connections; take the phototransducing property of cones (the property of releasing neurotransmitters in response to light) to be realized by, among other properties/relations, the light absorbing property of photopigment molecules; and take the property of absorbing light of a certain spectrum, of individual photopigment molecules, to be realized by,

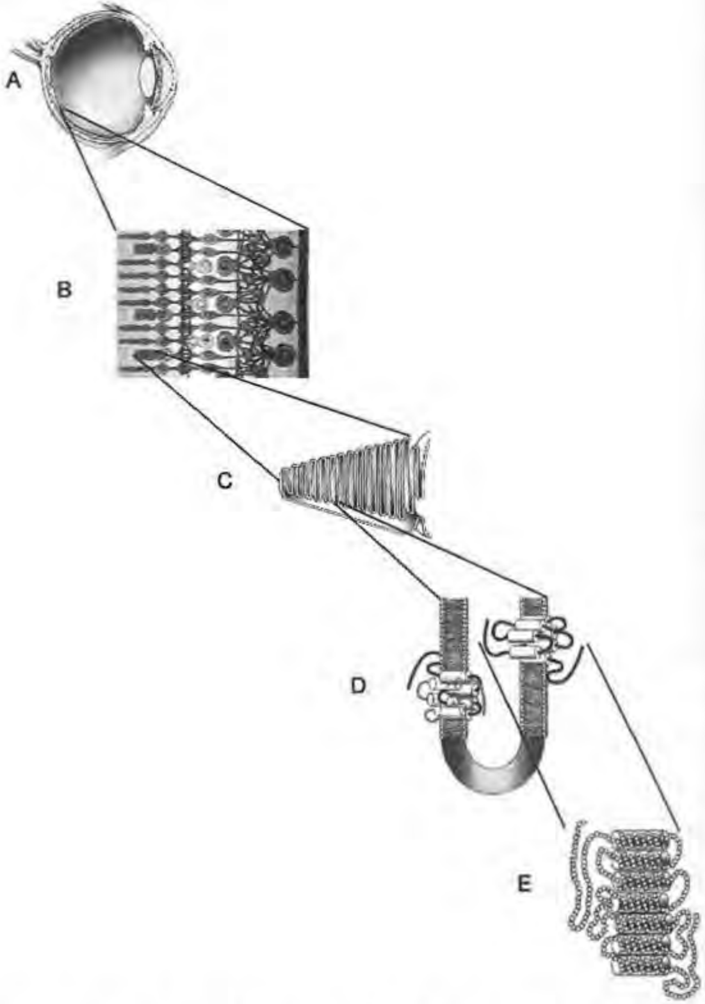

Figure 22.1. (A) The eye. (B) A Cross-section of the retina showing the principal cell types (including the rods and cones). (C) The outer segment of a cone. (D) Photopigments embedded in the membrane of the cone outer segment. (E) The amino acid chain of a cone photopigment.

among other properties/relations, the valence properties of the constituent atoms and their bonds.

Our approach is to work our way up through the mechanistic explanations offered at successively higher levels, starting with how atoms and their properties/relations constitute and realize photopigments and their properties—thus working from the bottom of figure 22.1 upward through the associated layers of explanations. By working through these various levels, we develop our various points in stages. First, we illustrate some general features of scientific composition, and then we articulate precise schemata for the realization and multiple realization of properties. In addition, our work also highlights the type of evidence that grounds the twin beliefs of neuroscientists that there are many neurobiological levels and individual variation at each of them. Last, and perhaps most important, we use our scientific examples to illuminate the reasons we contend that the evidence supporting these neuroscientific beliefs also underpins pervasive or "massive" multiple realization in neurobiology.

To start, let us consider an atomic-to-molecular case, where the lower level individuals are atoms of hydrogen, carbon, oxygen, nitrogen, and so on, and the higher level individual is a molecule of normal human green photopigment.[6] The relevant lower level properties and relations of the atoms include charge and bonding relations; the relevant higher level property of the photopigment is the property of being maximally sensitive to light of about 530 nm, with a bell-shaped distribution of sensitivity around that peak (see figure 22.2).[7] The sciences provide a clear mechanistic explanation of why a normal human green photopigment has the latter property under the normal physiological background conditions in a cone.

The sciences distinguish two portions of the cone photopigments: an 11-*cis*-retinal element and an opsin protein. The individual atoms in the retinal element have properties such as size and valence, which give them the powers to form certain types of bonds in response to various situations. For instance, the bonded and spatially aligned carbon, hydrogen, and oxygen atoms in a molecule of 11-*cis*-retinal form a long chain of alternating single and double carbon bonds (see figure 22.3). In this chain, the bond between the 11th and 12th carbon atoms of the 11-*cis*-retinal contributes the power of capturing a photon of light of a certain kind to these atoms. As a result, the powers contributed by the atoms' properties and relations noncausally result in the green photopigment having the property of absorbing a particular spectrum of light with a maximum sensitivity at 530 nm. The properties and relations of the individual atoms thus together realize the photopigment's property.

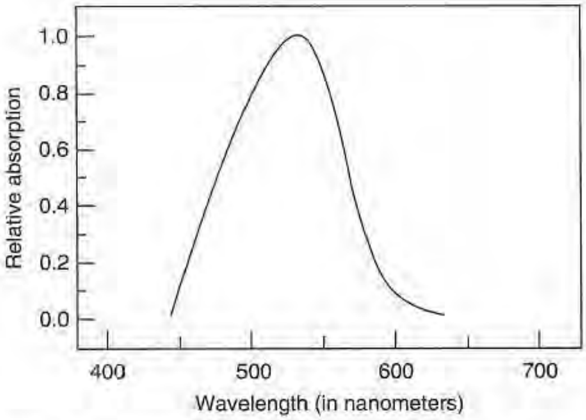

Figure 22.2. M-cone photopigment sensitivity curve. Modified from Sekular and Blake (2002, figure 2.23, p. 74).

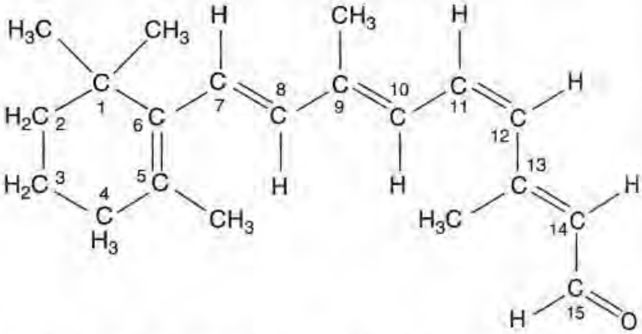

Figure 22.3. 11-cis-retinal. Modified from Casiday and Frey (1998).

In this and other mechanistic explanations, we have compositional relations posited between powers, properties, individuals, and processes. Individual atoms of carbon, nitrogen, hydrogen, and so forth constitute the photopigment molecule. The powers of the atoms to capture certain photons comprise the power of a photopigment molecule. The properties and relations of these individual atoms together realize the light sensitivity of the photopigment molecule and the processes of assimilating photons into a particular electronic configuration implement the process of absorbing light of a certain spectrum. Though such complexity is daunting, we can begin to understand such compositional notions if we use our example to draw out what appear to be some of their general features.

First, we note that the various compositional relations in our example are a species of determination relation, but one that is rather different from causal determination—such "horizontal" causal determination is temporally extended, relates wholly distinct entities, and often involves the transfer of energy or the mediation of force. In contrast, the vertical determination involved with compositional relations is synchronous. For example, it takes no time for the atoms to constitute the photopigment molecule or for the properties and relations of the constituent atoms to realize the property of absorbing light of 530 nm. Compositional relations also do not relate wholly distinct entities, because it is the individual atoms that constitute the photopigment molecule and the properties of the atoms (such as their size, charge, polarity, bonding relations, etc.) that realize the higher level property (such as a green photopigment molecule's light sensitivity). Finally, compositional relations do not involve the transfer of energy and/or the mediation of force between composing and composed entities. Compositional relations in the sciences are thus very different from causal relations and are a variety of what we call noncausal determination.

Second, compositional relations in the sciences usually relate qualitatively different kinds of entity. For example, the green photopigment has the property of being maximally sensitive to light of 530 nm, but no atom in the photopigment has such a property. We thus have individuals that constitute other individuals with which they need share no properties. Similar points hold for the relevant powers, properties, and mechanisms. A quick examination of cases of compositional relations posited in the special sciences shows that this feature is pervasive—entities usually compose entities of *qualitatively different* kinds.

Initially, it might seem surprising that entities of one kind could compose and explain entities of completely different kinds. Third, and perhaps most important for our purposes, we should mark that compositional relations are usually "many-one" in the sense that *many* component entities compose some higher level entity. Thus a number of atoms constitute the molecule of photopigment, and a number of properties and relations of the atoms realize the property of absorbing light of 530 nm. This feature is important because it dispels any mystery about how relations of composition in the sciences can relate qualitatively different entities. Even though the composing entities are individually different from the composed entity, nonetheless the composing entities *together* noncausally result in the composed entity. This distinctive feature of such composition relations consequently allows one to mechanistically explain powers, properties, individuals, and processes of one kind using, *together*, powers, properties, individuals, and processes of very different kinds.

As we will see shortly, the latter mundane feature also underlies the phenomenon of multiple realization, but before we turn to multiple realization, let us now more carefully articulate the nature of scientific realization. In our case, the photopigment's property of being maximally sensitive to light of 530 nm is individuated by the power of absorbing light in the neighborhood of 530 nm. As we outlined, the mechanistic explanation of why the photopigment has this property is that its constituent atoms have properties, such as valence and charge, which contribute the powers of capturing photons and changing their electronic configurations to form new sets of bonds. As a result, the sciences tell us that the photopigment molecule has a specific property of absorbing light of 530 nm, G, because its constituent atoms have properties and relations, F_1–F_n, that can change their energy levels in a very particular way on absorption of certain photons. The powers contributed by and individuative of the properties and relations of the constituent atoms in this manner noncausally result in the powers contributed by and individuative of the property of the photopigment.

Using these observations as a guide, we offer this thumbnail account of realization in the sciences (elsewhere dubbed the Dimensioned view):[8]

> (Realization) Property/relation instance(s) F_1–F_n realize an instance of a property G, in an individual *s* under conditions \$, *if and only if*, under \$, F_1–F_n together contribute powers, to *s* or *s*'s part(s)/constituent(s), in virtue of which *s* has powers that are individuative of an instance of G, but not vice versa.

A number of features of the Dimensioned view, mirroring the common characteristics of scientific composition noted earlier in our example, are worth emphasis. First, the Dimensioned view accommodates realization as a species of noncausal determination. Second, it permits realizer and realized properties to be qualitatively distinct, allowing that these properties may contribute no common powers and be instantiated in different individuals. Perhaps most important, the Dimensioned account implicitly acknowledges that realization is usually a many-one relation, for it allows that many realizer properties may contribute powers that *together* determine that the relevant individual has the qualitatively different powers individuative of the qualitatively different realized property.

Our work understanding the common features of composition in the sciences obviously underpins this view of realization and it again offers help if we turn to the phenomenon of multiple realization. Recall the second and third of the general features of scientific composition, which are shared by realization relations. Given the characteristic that scientific realization often relates qualitatively different kinds of property, and the feature that many properties together realize other properties, then a variety of realizer properties that are qualitatively distinct from other realizers and the realized property can each *together* realize instances of the same special science property. The result is *multiple* realization—instances of the same higher level property realized by distinct lower level properties and relations that together noncausally result in the powers of the same special science property, despite being different from each other and the realized property in the powers they individually contribute.

We can give substance to these abstract points if we again return to the concrete example of the green photopigment and its property of maximally absorbing light of 530 nm. For the sciences have now identified two chemically distinct molecules, constituted by two distinct combinations of atoms, that current evidence indicates have the same peak sensitivity as the normal human green photopigment (see, for example, Merbs and Nathans, 1993). In addition to the normal amino acid sequence of the green opsin, there is another sequence produced by a homologous recombination of the first two exons of the gene for the normal human red photopigment with the last four exons of the gene for the normal green photopigment. As a result, given the differing properties and relations of the atoms in these two molecules, there are two known combinations of atomic properties and relations that noncausally give rise to the same property of maximally absorbing light of 530 nm.[9] We thus have different realizations of the standard green peak light sensitivity. This should be unsurprising, for we have seen that because realizers usually compose a qualitatively different realized property, this opens the space for *distinct* combinations of realizers to noncausally result in instances of the *same* higher level property. In just this fashion, we have multiple realizations at the atomic level of the property of maximally absorbing light of 530 nm.

(As an aside, because this type of point will be important later, notice that the multiple realization of the property of being maximally sensitive to light of 530 nm is not simply a function of our attending to a property using a relatively coarse "grain" of description at the higher level, such as being light sensitive, rather than a relatively fine "grain" of description, such as being maximally sensitive to light of 530 nm. Even the relatively fine grain of description of the higher level property allows for its multiple realization. Anyone familiar with the recent literature will recognize that we are reacting to concerns raised by Bechtel and Mundale, 1999. In section 3, we return to Bechtel and Mundale's point about grains of description and directly address objections that may be based on their concerns.)

We can use these points to frame a precise, abstract account of multiple realization in the sciences as follows:

> (Multiple Realization) A property G is multiply realized *if and only if* (i) under condition $\$$, an individual *s* has an instance of property G in virtue of the powers contributed by instances of properties/relations F_1–F_n to *s*, or *s*'s constituents, but not vice versa; (ii) under condition $\$^*$ (which may or may not be identical to $\$$), an individual s^* (which may or may not be identical to *s*) has an instance of a property G in virtue of the powers contributed by instances of properties/relations F^*_1–F^*_m of s^* or s^*'s constituents, but not vice versa; (iii) F_1–$F_n \neq F^*_1$–F^*_m; and (iv), under conditions $\$$ and $\$^*$, F_1–F_n and F^*_1–F^*_m are at the same scientific level of properties.[10]

Overall, the theory schema is fairly obvious. Conditions (i)–(iii) simply frame the demand for distinct sets of realizer properties for instances of the same realized property. However, the final condition deserves more comment.

Implicitly, philosophers have always had something like condition (iv) in mind when discussing multiple realization in the sciences. To see why one needs (iv), either implicitly or explicitly, consider the following common situation. Properties and relations of certain atoms realize the property of maximally absorbing a certain frequency of light; but obviously properties and relations of certain fundamental microphysical properties realize the properties and relations of these atoms and hence *also* realize this instance of the property of the photopigment of maximally absorbing that frequency of light. But since the properties and relations of the atoms ≠ properties and relations of fundamental microphysical individuals, it appears that in such cases if we only use conditions (i)–(iii), then this entails we have a case of multiple realization. But we obviously do not want to treat the difference between realizers at the physical and chemical levels as sufficient for multiple realization. What has gone awry is that the two sets of properties are not at the same level and are implicitly excluded as candidates to ground a case of multiple realization. Addition of condition (iv) explicitly resolves this problem, though we suggest the condition is usually implicitly accepted as a shared background condition in earlier discussions of multiple realization in the sciences.[11]

An advantage of using (iv) is that it also combats a common philosophical practice that can cause problems. The practice in question is that of talking simply about the multiple realization of some property, whether psychological, biological, or whatever, and saying nothing further. Often, given the context, this may be a harmless way of talking, but we should note how it may be damaging. Given the nature of the realization relation, claims about realization and hence multiple realization are always relative to *particular* properties and levels—as both of our schemata now make explicit. Thus, property instance G is not simply realized; rather, it is realized by instances of certain lower level properties F_1–F_n. And instances of property G are multiply realized by instances of properties F_1–F_n and instances of properties F^*_1–F^*_m, when F_1–F_n and F^*_1–F^*_m are at the same level, as condition (iv) now makes clear. Thus, claims of realization and multiple realization are always indexed to particular levels and specific properties at these levels. We can quickly see the importance of this point.

Suppose that some higher level property G is multiply realized by microphysical properties of fundamental particles and hence multiply realized at the microphysical level. This does not, of course, mean that G is multiply realized in, say, distinct physiological properties. After all, it is logically possible to have G be univocally realized in the same physiological properties of two organisms and also have these properties in turn be univocally realized in the same biochemical properties of these organisms, but then have these biochemical properties be multiply realized by the microphysical properties of the fundamental particles that constitute these two organisms. So our two instances of G might be univocally realized at level X (the physiological level) and level $X - 1$ (the biochemical level), and still be multiply realized at level $X - 3$ (the microphysical level). We can thus see that a property is not simply *either* univocally realized *or* multiply realized. This is a false dichotomy, for such ascriptions are indexed to levels, and a property may be univocally realized at one level and multiply realized at another. To avoid confusion in talking about multiple realization, one therefore needs to be careful to make claims about realization, and hence multiple realization, indexed to particular realizers and levels.

If we return to our general accounts of realization and multiple realization, we can further illustrate their character if we consider another layer of mechanistic explanation we find for color processing in the retina. We have already noted how properties and relations of atoms can realize and multiply realize a molecular property. So let us move to a molecular-to-cellular case in our molecular explanations of the properties of a human cone at the cellular level, where we again consider how molecular properties and relations realize and multiply realize a cellular property.

In this case, at the lower level the relevant individuals are water molecules, ions (such as Ka^+, Na^+, and Ca^{++}), phospholipids, proteins, and so on, and at the higher level the individual under consideration is obviously a human cone (see figures 22.4 and 22.5). The higher level property of the cone that is mechanistically explained in this case is its property of releasing a neurotransmitter, in this case glutamate, in response to the absorption of light. The lower level properties and relations used to explain this property include having a charge, light sensitivity, polarity, and spatial arrangement.

In this case, our mechanistic explanations are more complex, but consider some of the highlights of these accounts of how the lower level entities compose (and hence explain) the higher level entities in question. Phospholipid molecules have both a hydrophilic and a hydrophobic region. Given this configuration, they

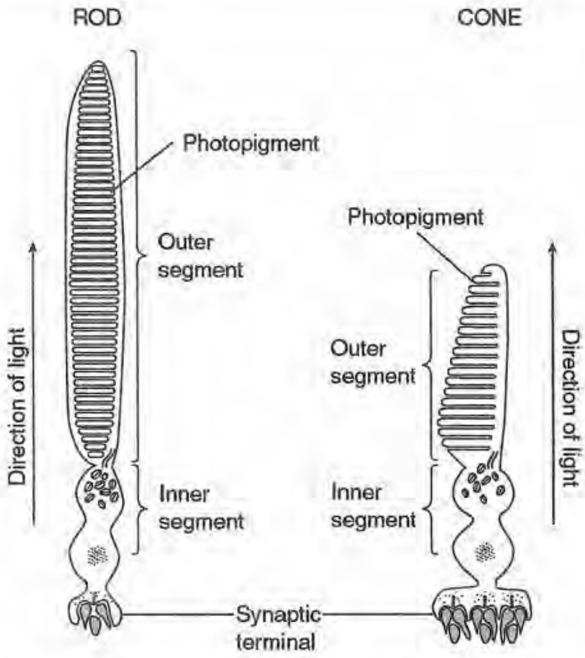

Figure 22.4. **Structure of human rods and cones. From Sekular and Blake (2002), figure 2.29, p. 69.**

spontaneously form a bilayer structure in which the hydrophilic regions face outward to an external aqueous environment in either the extracellular space or the cytoplasm, while the hydrophobic tails of the molecules cluster together inside the bilayer. This phospholipid bilayer constitutes the cell membrane, illustrated in the right half of figure 22.5. Proteins, for their part, also have hydrophobic and hydrophilic portions that help embed them in the cell membrane (see again the right half of figure 22.5). Human cone opsins, for example, have an evolutionarily well-conserved set of seven transmembrance amino acid sequences (see figure 22.6). Ion channels have amino acid sequences that enable them to span the cell membrane and provide bindings sites on one or another side to regulate the flow of ions through the channel. Cytoskeletal proteins, also partially embedded in the cell membrane, shape a cell into exotic configurations, such as those of the rods or cones.

Photopigment molecules are embedded in the cell membrane in the outer segment of the cone (recall figure 22.5). On absorption of a photon, a single photopigment molecule will change conformation and release into the cytoplasm a molecule of all-*trans*-retinal leaving an activated opsin molecule in the membrane. One activated opsin molecule binds to a single G protein molecule located on the inner surface of the cell membrane. This G protein molecule, in turn, activates a molecule of an enzyme, cGMP phosphodiesterase, which breaks down cGMP. When

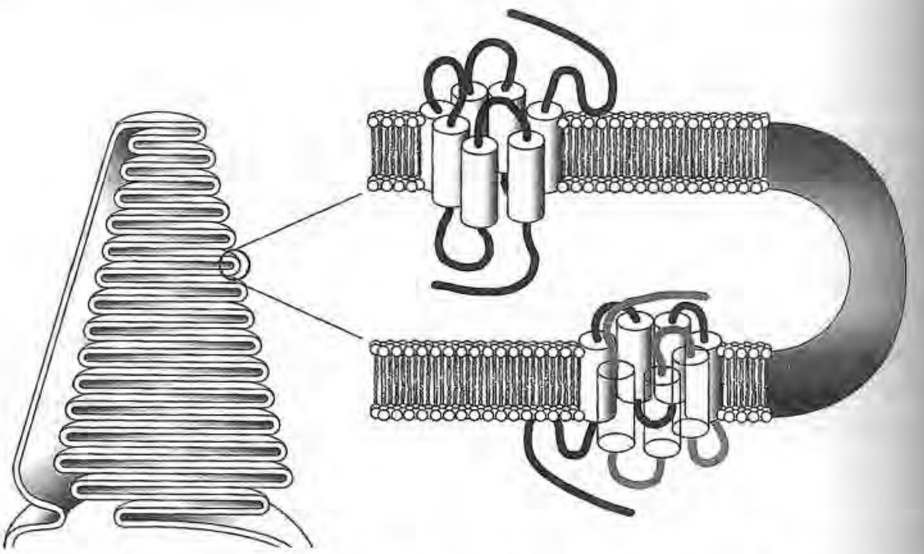

Figure 22.5. Photopigment molecule embedded in the cell membrane and phospholipid molecules of the membrane constituting a cone. Modified from Sharpe, Stockman, Jägle, and Nathans (1999), figure 1.2, p. 6.

intracellular cGMP concentrations subsequently decrease, cGMP is removed from a cGMP-gated Na$^+$ channel, leading to the closure of the channel. Closing the channel blocks the influx of Na$^+$ into the cell. In concert, vast numbers of photopigment molecules, G protein molecules, ion channels, and Na$^+$ ions go through this process, leading to the hyperpolarization of the cell. This hyperpolarization propagates from the outer segment to the synaptic contact of the cone, where it reduces the rate

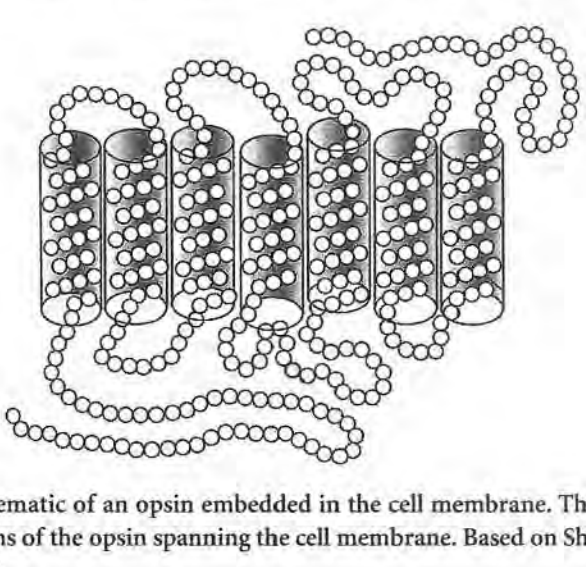

Figure 22.6. Schematic of an opsin embedded in the cell membrane. The seven cylinders represent portions of the opsin spanning the cell membrane. Based on Sharpe et al. (1999), figure 1.17A, p. 43.

of release of the neurotransmitter glutamate. This reduction in neurotransmitter release is the cone's signal that the cell has been illuminated. The foregoing lower level processes may be summarized schematically as follows.

Photon capture → all-*trans*-retinal release → G protein activation →
cGMP phosphodiesterase activation → cGMP decrease →
cGMP released from ion channels → ion channel closure →
cone hyperpolarization → decreased glutamate release.

Obviously a large number of these molecular processes occur together, and these lower level processes implement the cellular process of signaling the presence of light by release of glutamate. Consequently, we can thus also see that the cone's property of releasing a neurotransmitter in the presence of light is evidently realized by the properties and relations of the molecular individuals within the cell.

Our molecular-cellular example illustrates exactly the same features of the realization relation we described in the atomic-to-molecular example. First, the lower level properties and relations of the molecules stand in a synchronous, noncausal determination relation to the higher level property of releasing a neurotransmitter in the presence of light. There is no transmission of energy or mediation of force between the lower level properties and relations and the higher level property, where these properties are also not wholly distinct. Second, the relata in this realization relation are once more qualitatively distinct. The relevant determining properties and relations of the molecules are their charges, polarity, and light-absorbing capacity, where the determined property of the cell is its releasing glutamate in response to the presence of light. (With regard to individuals, the particular molecules in a cone do not release glutamate in response to light, whereas the cone does have this property. Similar points hold for the relevant higher and lower level powers and processes.) Third, the property of releasing a neurotransmitter in response to light is realized by *many* molecular properties and relations. It is the properties and relations of the individual molecules that *together* result in the cell's property of releasing glutamate in response to illumination.

The foregoing explains how a cone's property of signaling is realized by the lower level properties of ions, phospholipids, proteins, and so on. Now, however, we can see how distinct sets of molecular-level properties can provide for multiple realizations of a cone's property of transducing light into glutamate release. To do this, we might focus on any of the differing relevant properties of any of the different protein molecules in the biochemical cascade already described. We might focus on the different molecules of cGMP phosphodiesterase. Still, the clearest case of multiple realization emerges from the research on the most studied components in the cascade, namely, the photopigments. These photopigments differ in one of their molecular level properties, namely, their absorption spectra.[12] We can thus see in this case that there are distinct molecular-level properties, that is, distinct absorption spectra, that give rise to the same cellular property of transducing light into a neurochemical currency of glutamate release. Once again, we have a case of multiple realization, in this example of a cellular property by molecular properties.

At this point, it is worth stepping back from our examination of concepts of scientific composition to note some related methodological practices. Working scientists move freely between appealing to entities at different levels, for instance, switching from focusing on atomic-level properties, in the differences in amino acid sequences, to appealing to molecular-level properties, in the differences in absorption spectra. In fact, working scientists often freely move up or down through a number of levels of compositionally related entities to get explanatory help. What underpins this amazingly fruitful methodological maneuver? Briefly put, researchers pursue this methodological strategy because they recognize that when entities bear compositional relations, the nature of component entities at one level has ramifications for the entities at some other level that they compose. For example, working scientists clearly appreciate the potential ramifications of differences in atomic-level properties for the molecular-level properties they are taken to realize.

It is worth marking that our schemata for realization and multiple realization provide a ready explanation for this common feature of actual science. Although we have not explicitly noted this feature so far, we should mark that our schemata take realization to be a *transitive* relation—a feature that it shares with other scientific composition relations. Thus if property instances F_1–F_m realize G_1–G_n, in certain individuals and under specific conditions, and G_1–G_n realize property instance H, then F_1–F_m realize H. Consequently, under our schema for scientific realization, it makes as much sense to say that an instance of the property of releasing a neurotransmitter in response to illumination (by light in the neighborhood of 530 nm) is realized by molecular-level properties as it does to say that this instance of the property of releasing glutamate in response to illumination (by light in the neighborhood of 530 nm) is realized by certain atomic level properties.[13] Given the transitivity of compositional relations like realization, we can thus see how working scientists can successfully move up or down through levels of compositionally related entities to provide explanatory gain in their work at some other level of entities. In section 3, we explore this important point further, but at the risk of belaboring the case of color processing in the retina, we wish to advance to yet one more layer of mechanistic explanation and another level of entities.

The last case we consider is a cell-to-tissue case that concerns how a certain property of the retina, basically the property of signaling a pattern of color in the visual field, is realized and multiply realized by cellular level properties. In this example, the lower level individuals are the particular photoreceptor cells, the amacrine cells, bipolar cells, horizontal cells, and retinal ganglion cells (see figure 22.7). The higher level individual is a retina. The lower level properties and relations of the cells include releasing glutamate in the presence of light within a given band of frequencies, releasing certain neurotransmitters, binding certain neurotransmitters, having certain electrochemical synapses, and certain patterns of connectivity. The higher level property of the retina that is mechanistically explained is the retina's property of signaling a pattern of color in the visual field.

Once again, we have mechanistic explanations of the relevant properties that are highly detailed and rather complex in nature, so we only briefly review some of the

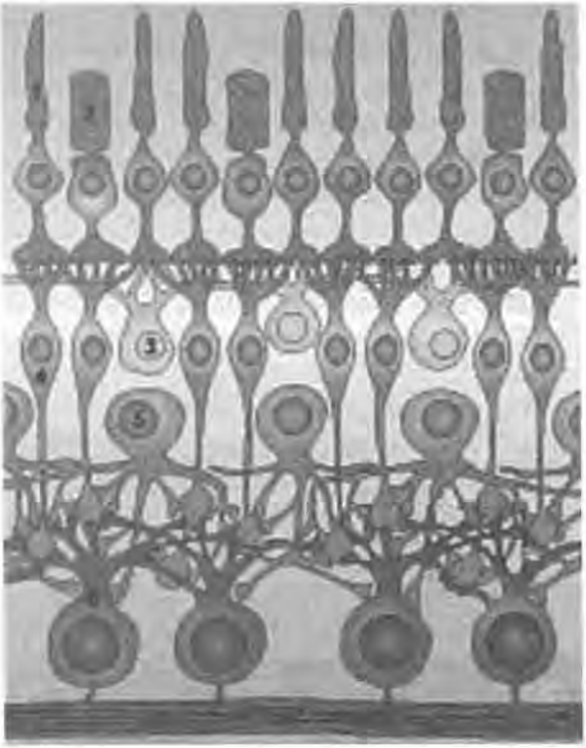

Figure 22.7. Cells constituting the human retina. (1) Rods, (2) cones, (3) horizontal cells, (4) bipolar cells, (5) amacrine cells, and (6) retinal ganglion cells. Modified from Wässle (2004), figure 1, p. 2.

highlights of our explanations of the retina's property in terms of the properties and relations of the cells that are taken to compose it. When introducing our molecular-cellular-level example, we already briefly described the nature of the biochemical cascade involved in translating photon capture into changes in glutamate release. Given only a single photopigment, a single cone can release glutamate in response to a relatively narrow band of light frequencies, but it is unable to signal the specific frequency of the incoming light. A given decrease in glutamate release may equally result from either a high-intensity light at a frequency to which the photopigment is relatively insensitive, or a low-intensity light at a frequency to which the photopigment is relatively sensitive. If processes of glutamate release from cones are going to implement a retina's process of signaling distinct patterns of color in the visual field, there is an obvious problem. However, this difficulty is resolved by the ratio of glutamate release in cells containing photopigments of different sensitivities. The three types of cones in the normal human eye, S-, M-, and L-cones, process short-, medium-, and long-wavelength frequencies of light. That is, each changes its glutamate release in response to a different band of frequencies of incident light (see figure 22.8). Each type of cone releases glutamate as it does in virtue of containing a chemically distinct photopigment. That is, each photopigment consists of a protein component, an opsin,

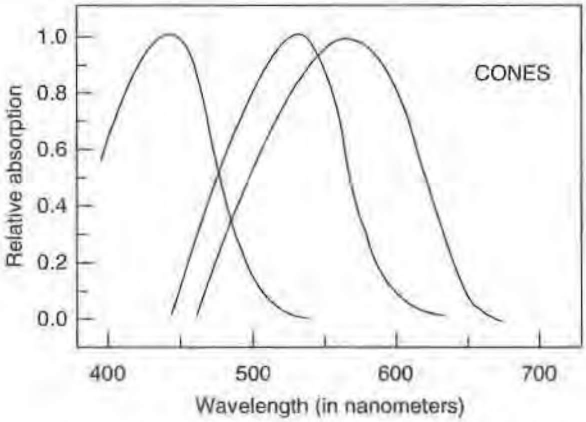

Figure 22.8. Absorption spectra of the S-, M-, and L-cones. From Sekular and Blake (2002), figure 2.23, p. 74.

covalently bonded to an 11-*cis*-retinal component, and the opsin components vary from cone to cone. The amacrine, bipolar, horizontal, and ganglion cells, of course, contribute to the retina as well, but for simplicity we set their role to one side.

This cell-to-tissue example again supports the existence of the features of the realization relation we described in our previous cases. First, the lower level properties and relations of the cells stand in a synchronous, noncausal determination relation to the higher level property of signaling patterns of color in the visual field. There is no transmission of energy or mediation of force between the lower level properties and relations of the cells to the higher level property of the retina, where the relevant properties, and the individuals that have them, are not wholly distinct. Second, the relata in this case of realization are again qualitatively distinct. The relevant determining properties and relations of the cells include their capacity to release glutamate in response to illumination, releasing certain neurotransmitters, binding certain neurotransmitters, having certain electrochemical synapses, having certain patterns of connectivity, and so on; in contrast, the determined property of the retina is its signaling patterns of color in the visual field. Third, in the case of the retina, once again *many* properties and relations of individual cells *together* noncausally result in the retina's property of signaling a pattern of color in the visual field.

There are many ways the property of signaling a pattern of color in the visual field is multiply realized by properties at the cellular level. The simplest examples stem from some of the principal forms of color blindness in which one type of cone is missing. The retinas of so-called dichromats are not completely insensitive to color; it is not as though they can make no color discriminations. Instead, they are able to make fewer color discriminations than can the retinas of normal humans. So these retinas are still color processors. There are, however, three ways of being a dichromatic retina, each corresponding to the loss of a distinct photopigment. The retina of a protanope lacks red cones, the retina of a deuteranope lacks green cones, and the retina of a tritanope lacks blue cones. Each form of dichromacy corresponds to a distinct realization of an instance of the property of signaling patterns of color in the visual field.

Less dramatic, and perhaps somewhat less familiar, are the cases of what are called anomalous trichromats. These individuals possess three distinct types of cones, but their sensitivities are not those of normal cones. For example, in one of the most common forms of red-green colorblindness, the red cones and the green cones release glutamate in response to relatively similar bands of electromagnetic radiation, hence there is not enough difference in the properties of the red and green cones to implement a higher level process that can signal differences between certain patterns of color. This gives rise to color-sensitivity "blind spots." Anomalous trichromacy is, of course, a kind of color processing, and its subjects do realize instances of the property of signaling patterns of color in the visual field, but in addition it grades off into normal color vision. There is no sharp dividing line between normal trichromacy and anomalous trichromacy. To take just one well-known example in the biochemistry of vision literature, the human red cone appears to be polymorphic. That is, it comes in two forms. One form of the red cone has the photopigment with an amino acid sequence with a serine amino acid at position 180, and the other photopigment has alanine at that position. Both forms are quite common in the human population. One has the multiple realization of normal human color vision by having some individuals with cones containing one red photopigment and other individuals with cones containing the other.[14]

Recent work on the biochemistry of opsins has suggested an even more radical form of multiple realization of the property of signaling patterns of color. In a classic paper, Nathans, Thomas, and Hogness (1986) identified the gene sequences, hence the amino acid sequences, of the three opsin components of the photopigments. In addition, they found that normal humans vary in the number of genes coding for the green pigment. In other words, there are multiple loci each coding for a green photopigment. Subsequent research has also found that normal humans vary in the number of genes coding for the red pigment (Neitz and Neitz, 1995). This suggests the possibility that a given individual can possess distinct versions of the gene for the green and red photopigments at the different loci, and hence can possess distinct green and red photopigments and distinct green and red cones. Furthermore, it is hypothesized that part of the reason for individual variation in color sensitivity within humans is due to differences in the number of different kinds of cones. Some individuals might have, say, only one type of green cone and one type of red cone, and other individuals might have, say, two different green cones and seven different types of red cone.[15]

To summarize, by examining a number of connected examples of mechanistic explanations from a familiar and well-confirmed area of research in neurobiology, we have supported a number of claims. First, we have shown that the compositional relations posited in such explanations have some important common features in being transitive, noncausal determination relations that are usually many-one and relate qualitatively different entities at distinct levels. Second, building on these observations about the common features of scientific composition generally, we have provided precise accounts of both scientific realization and multiple realization. As we have seen, once one gets a better grip on scientific composition, the character of multiple realization becomes more understandable. For we have seen that a number of component entities usually together compose some qualitatively different entity. As a result, diverse kinds of component entities can together, usually on separate occasions, compose instances of the same composed entity. The result, in the case of properties, is multiple realization of instances of the same higher level property by different lower level realizers.

In addition to articulating a theoretical framework for realization and multiple realization, our work surveying these concrete scientific examples also supported our specific claims about neurobiology. In the case of color processing in the human retina, we have shown that there are a number of neurobiological levels relevant to our mechanistic explanations. This supports the first belief we have attributed to working neurobiologists. Furthermore, at the atomic-to-molecular, the molecular-to-cellular, and cell-to-tissue cases, we have also shown that we find individual variation among the relevant entities at the lower level, thus supporting the second belief of neurobiologists that we find individual variation at all levels. Finally, applying the theoretical framework to the empirical evidence that underlies scientists' commitments to multiple levels and individual variation in each of them, we have also shown how a strong case can be made for multiple realization of molecular properties by properties at the atomic level, cellular properties by properties at the molecular level, and tissue properties by properties at the cellular level. Given the transitivity of scientific relations of realization, in our concrete scientific case we have thus found plausible evidence for the multiple realization at many neurobiological levels of the relevant human psychological properties in this example, such as the property of signaling patterns of color in the visual field.

2. Levels, Individual Variation, and Multiple Realization in Nervous Systems

With our theoretical framework in place, we now wish to widen our focus and examine evidence about a still wider range of entities in various neurobiological levels. Obviously we cannot attempt anything like a comprehensive survey of the relevant empirical findings; in fact, we can only highlight a fraction of these results. However, our initial goals will be to note the apparently wide body of evidence supporting the twin claims of working neuroscientists: that there are a number of neuroscientific levels and that individuals at each of these levels vary in their properties. Our work in the previous section highlights why such findings plausibly underpin multiple realization, so our approach is to survey the evidence for individual variation at each level and then conclude the section by returning to the issue of why such evidence provides support for multiple realization. Given the recent and rapidly growing resistance of so many philosophers to the existence of multiple realization in psychology and neuroscience, we hope that our work, however rough and incomplete, provides a clear and bold statement of the type of evidence that such critics need to address.

We begin, for obvious reasons, with lower neuroscientific levels, where our accounts are more mature and work through progressively higher levels, where our understanding becomes steadily less developed. As the reader will clearly see, a pattern of individual variation is obvious across these levels. Given our focus on multiple realization and the relative stages of development of various areas of research, our focus in each case is on showing that at the level in question there is variation across individuals where, very roughly, we have the following situations. Either (i) the properties of one level are known to be realizers of properties at the next higher neuroscientific level; (ii) our present accounts suggest that the properties of one level are likely to be realizers of properties at the next higher neuroscientific level; or (iii) our accounts of the properties of one level do not yet enable us determine which specific properties at the lower level are likely to be realizers of properties at the next higher neuroscientific level. Our focus on (i), (ii), or (iii) obviously results from the very different stages of development of both our accounts of the entities within different levels of neuroscientific organization and/or our intralevel mechanistic explanations of entities at differing levels using compositional relations between these entities. In concluding the section, we argue that the pattern of pervasive individual variation, in situations (i), (ii), or (iii), nonetheless grounds a plausible case that the relevant neuroscientific properties are multiply realized.

In section 1, we reviewed the scientific case for the view that there is multiple realization by properties at the atomic level of a molecule of photopigment's property of having a particular light sensitivity. Now we wish to draw attention to the fact that the properties of human cone photopigments are unlikely to be unique in this regard and that many molecular properties are likely to be similarly multiply realized at the atomic level. Consider a hemoglobin molecule's property of binding oxygen with such and such an affinity, or a protein kinase A (PKA) molecule's property of phosphorylating cAMP response element binding protein (CREB) with such and such a rate constant, or an alcohol dehydrogenase molecule's property of oxidizing ethanol to acetaldehyde at such and such a rate. Each of the foregoing molecular properties is likely to be realized by distinct combinations of atomic level properties. Thus, one instance of the property G of binding oxygen with such and such an affinity will be realized by one set of atoms bearing properties/relation F_1–F_n (such as valence and bonding), where another instance of G will be realized by another set of atoms bearing properties/relations F^*_1–F^*_m, where F_1–F_n and F^*_1–F^*_m are at the same level, and F_1–$F_n \neq F^*_1$–F^*_m. Similar points hold mutatis mutandis for other properties such as G_1 of phosphorylating CREB.[16]

We can quickly identify some of the reasons for such widespread multiple realization at the atomic level. Any given type of protein varies in its amino acid sequence across different individual organisms, so that any given type of protein varies in the numbers and arrangements of atoms of carbon, hydrogen, nitrogen, oxygen, and so on. Along with this variation in individuals comes variation in their properties. Atoms of carbon, hydrogen, nitrogen, oxygen, and so forth differ in (most notably) their valence, which confers on them different powers to bind to other atoms. The properties/relations of the atoms give proteins such properties as their size, shape, and charge distributions. Properties like size, shape, and charge distribution confer on proteins their powers to bind substrates and catalyze biochemical reactions. Finally, along with this variation in properties/relations comes variation in powers. The different binding powers of a protein's atoms comprise the different powers for such things as binding substrates and catalyzing biochemical reactions.[17] The foregoing is a more general outline of the underlying factors we found for the cone photopigments examined in section 2, and we have now seen that such factors appear to apply quite generally for properties of complex proteins at the biochemical level. We have thus found strong evidence for both individual variation at the biochemical level and multiple realization of the properties at this level by atomic properties, thus indicating this is an example of type (i).

Move, now, to a slightly higher level of neurobiological organization, where we find dendritic spines. Dendritic spines are individual finger-like to mushroom-like extensions on the dendrites of neurons to which synapses connect (see figure 22.9). Individual dendritic spines are constituted by various types of individuals, including phospholipid molecules, water molecules, various individual ions, various cytoskeletal protein molecules that support the spine's shape and function, and various proteins embedded in or attached to the membrane surface. Properties such as a given size and shape of an individual dendritic spine are realized by various properties of the spine's constituents. The individual constituents of a dendritic spine bear such properties as having a hydrophobic segment (the phospholipids),

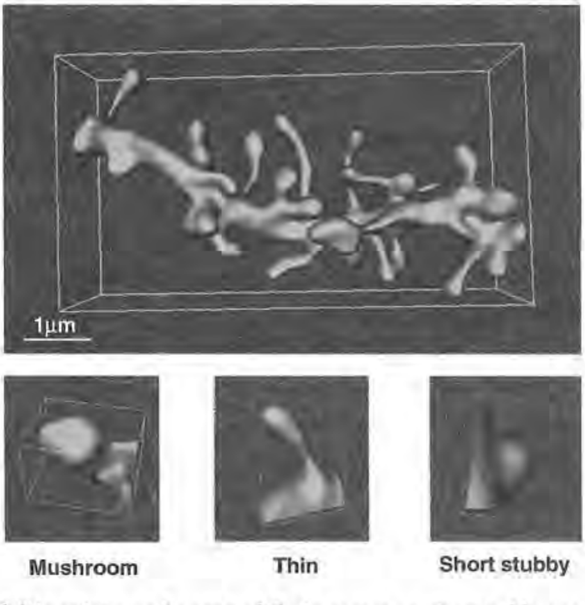

Figure 22.9. Dendritic spines and some of their common shapes. No synapses from other neurons are shown. McKinney (2005), figure 1, p. 1300. Reproduced with permission.

having a hydrophilic segment (the phospholipids), being polar (water), bearing a positive charge (K^+, Na^+, Ca^{++} ions), gating ions (K^+ and Na^+ channels), and so on. Most relevant for determining, say, the shape of the spine might be certain cytoskeletal proteins. These cytoskeletal proteins multiply realize the shape of a dendritic spine insofar as different protein molecules contribute different properties that noncausally determine the shape of the spine. Applying our schemata for realization and multiple realization, we would say that in some cases it is the instances of properties/relations F_1–F_n of the individual proteins of a dendritic spine that together realize some of the instances of the properties G of having a mushroom shape, where in other cases it is the instances of properties/relations F^*_1–F^*_m of the individual protein molecules of the spine that realize other instances of G.[18]

We have now seen that the properties of a dendritic spine might be multiply realized by lower level properties and that we presently have a case of type (ii) with such properties. But at the still higher level neuronal level of properties, dendritic spines are still more interesting for their apparent role in multiply realizing properties of neurons. An individual neuron is, of course, constituted by numerous subcellular components, such as dendritic spines, dendrites, axons, Golgi apparati, endoplasmic reticuli, and so forth. These individuals bear all manner of properties, such as size and shape, and we might, in principle, look to the properties of any of these constituents for sources of multiple realization of a neuronal property G, such as a V1 neuron's property of responding maximally to a line of a particular orientation. To illustrate our points, however, we focus on the properties of individual dendritic spines to illuminate sources of multiple realization at the neuronal or cellular level.

Dendritic spines differ in size and shape as measured in terms of such things as the size of the neck, their length, total volume, and head volume (see table 22.2 for data on hippocampal dendritic spines). Differences in these properties confer different electrical powers on these spines. For example, a spine's property of having a long neck and large head confers on it the power to limit the effects of the neurotransmitter released onto that spine to that spine. Furthermore, a spine's property of having a long neck and large head confers on it the power to limit any firing-dependent chemical changes in the synapse to that spine. If we consider the properties of all the dendritic spines on a given neuron, one can expect them to be such that there is one set F_1–F_n that will in part

Table 22.2. Variation in Functionally Significant Properties of Dendritic Spines: Ranges in Dimensions of Hippocampal Dendritic Spines and Their Synapses

	Dentate Gyrus	Area CA3	Area CA1
Neck diameter (μm)	0.09–0.54	0.20–1.00	0.038–0.46
Spine length (μm)	0.20–1.78	0.60–6.50	0.160–2.13
Spine volume (μm^3)	0.003–0.23	0.13–1.83	0.004–0.56
Head volume (μm^3)			0.003–0.55

Source: Sorra and Harris, 2000.

realize, say, a V1 neuron's property G of responding maximally to a line of a particular orientation, where another set of properties F^*_1–F^*_m of all the dendritic spines on another neuron will also realize that neuron's instance of property G.[19] Recent work on dendritic spines supports this view. Many of the basic features of dendritic spine morphology have been revealed by electron microscopic investigations of serial sections of fixed tissues, but more recent work has made it possible to make in vivo observations of changes in spine morphology over periods of days and weeks.

Scientists can insert genes that code for fluorescent proteins into mice. These genes produce fluorescent proteins in neurons that can be reliably reidentified day after day through fluorescence microscopy. With this technique, a large majority of dendritic spines have been found to be stable over the course of days and weeks, but nevertheless there remain, even in adult mice, changes in the number, size, and shape of dendritic spines.[20] In other words, this technique has revealed that even in adults there remains some degree of plasticity in spine morphology. Insofar as spine morphology varies in a single individual over time, it is highly likely that there will be variation in spine morphologies in different individuals of a given species at a given time.

Here we maintain that we have an example of type (ii), wherein the properties of one level are likely to be realizers of properties at the next higher neuroscientific level. What remains open for future investigation is the extent to which distinct properties/relations of distinct numbers and configurations of dendritic spines can be said to realize, at least in part, different properties of neurons.

Next consider a series of larger neural structures, where we think we can be less confident about which lower level properties realize which higher level properties. Thus we do not provide the level of detail given in earlier examples in this section where we plausibly had cases of types (i) or (ii). For example, consider cortical columns. Each individual column consists of a set of neurons that are relatively densely connected among themselves, but relatively less densely connected to neurons outside the column.[21] Many of the properties of individual columns, thus, appear to be realized at least in part by the properties/relations of individual neurons. Consider the finding by Kaschube, Wolf, Geisel, and Löwel (2002), using radiographic imaging techniques, that there is a high degree of individual variation in the size and shape of orientation columns in a population of 31 animals, some of which came from the same litter. Here we have a case of type (iii), but nonetheless we have strong evidence for individual variation at the level and the likelihood that the properties that are realizers will also be varied, hence making multiple realization likely.

Continue to another higher level of organization. In area V1, cortical columns are organized into still larger structures, ocular dominance columns. Ocular dominance columns are regions of layer IVc of area V1 that respond preferentially to inputs from one or another eye. Cytochrome oxidase staining enables the columns associated with one eye to be stained and the columns associated with the other remained unstained. In a study of six macaques, Horton and Hocking (1996) found numerous dimensions of variability in the ocular dominance columns. They found that the number, or periodicity, of ocular dominance columns in V1, varies by something on the order of 50 percent, and this variation is independent of the surface area of V1. Increased periodicity is, however, correlated with the complexity of the columnar mosaic. That realize, say, a V1 neuron's property G of responding maximally to a line of a particular orientation, where another set of properties F^*_1–F^*_m of all the dendritic spines on another neuron will also realize that neuron's instance of property G.[19] Recent work on dendritic spines supports this view. Many of the basic features of dendritic spine morphology have been revealed by electron microscopic investigations of serial sections of fixed tissues, but more recent work has made it possible to make in vivo observations of changes in spine morphology over periods of days and weeks.

Scientists can insert genes that code for fluorescent proteins into mice. These genes produce fluorescent proteins in neurons that can be reliably reidentified day after day through fluorescence microscopy. With this technique, a large majority of dendritic spines have been found to be stable over the course of days and weeks, but nevertheless there remain, even in adult mice, changes in the number, size, and shape of dendritic spines.[20] In other words, this technique has revealed that even in adults there remains some degree of plasticity in spine morphology. Insofar as spine morphology varies in a single individual over time, it is highly likely that there will be variation in spine morphologies in different individuals of a given species at a given time.

Here we maintain that we have an example of type (ii), wherein the properties of one level are likely to be realizers of properties at the next higher neuroscientific level. What remains open for future investigation is the extent to which distinct properties/relations of distinct numbers and configurations of dendritic spines can be said to realize, at least in part, different properties of neurons.

Next consider a series of larger neural structures, where we think we can be less confident about which lower level properties realize which higher level properties. Thus we do not provide the level of detail given in earlier examples in this section where we plausibly had cases of types (i) or (ii). For example, consider cortical columns. Each individual column consists of a set of neurons that are relatively densely connected among themselves, but relatively less densely connected to neurons outside the column.[21] Many of the properties of individual columns, thus, appear to be realized at least in part by the properties/relations of individual neurons. Consider the finding by Kaschube, Wolf, Geisel, and Löwel (2002), using radiographic imaging techniques, that there is a high degree of individual variation in the size and shape of orientation columns in a population of 31 animals, some of which came from the same litter. Here we have a case of type (iii), but nonetheless we have strong evidence for individual variation at the level and the likelihood that the properties that are realizers will also be varied, hence making multiple realization likely.

Continue to another higher level of organization. In area V1, cortical columns are organized into still larger structures, ocular dominance columns. Ocular dominance columns are regions of layer IVc of area V1 that respond preferentially to inputs from one or another eye. Cytochrome oxidase staining enables the columns associated with one eye to be stained and the columns associated with the other remained unstained. In a study of six macaques, Horton and Hocking (1996) found numerous dimensions of variability in the ocular dominance columns. They found that the number, or periodicity, of ocular dominance columns in V1, varies by something on the order of 50 percent, and this variation is independent of the surface area of V1. Increased periodicity is, however, correlated with the complexity of the columnar mosaic. That is, monkeys with relatively few columns have longer and smoother columns, where monkeys with relatively more columns also had shorter columns with more frequent bifurcations and islands (see figure 22.10). Using radiographic techniques, Kaschube et al. (2003) quantified the interindividual variability in the spacing of ocular dominance columns of cat primary visual cortex in 39 animals from three colonies. Once again, we have powerful evidence for individual variation in the properties at the relevant levels despite this being a case of type (iii) where we do not know which specific properties at this level are likely to be the realizers of particular higher level properties.

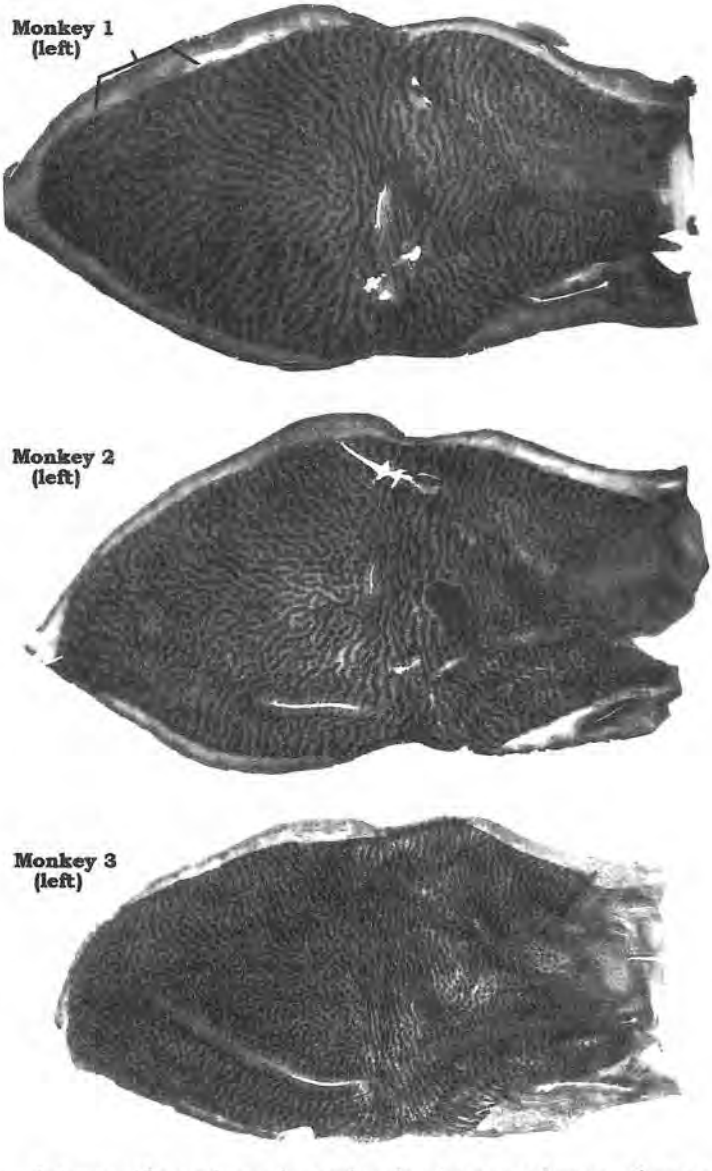

Figure 22.10. Intraspecific variance in ocular dominance columns of macaques. From Horton and Hocking (1996), figure 2, p. 7232.

Next consider some of the larger brain regions and collections of brain regions so convincingly described by Korbinian Brodmann (1909). Over the past century, a variety of methods have revealed individual variations in these structures made up of cortical columns. Stensaas, Eddington, and Dobelle (1974) found the human area V1 to vary in surface area by a factor of up to three.[22] Furthermore, some of these variations are highly correlated. Using structural and functional MRI, Dougherty et al. (2003) found that when V1, V2, and V3 are functionally delimited, the central 12° of the visual fields of V1 and V2 vary in surface area by a factor of about 2.5, where the central 12° of the visual field of V3 displays somewhat more limited variability. In addition, they found correlations in the variation of the surface area of V1 and V2 but weaker correlations between V1/V2 and V3. In primary auditory cortex, Rademacher et al. (2001) found that like V1, regions Te1 and Heschel's gyrus could vary by a factor of up to three across individuals. Penfield and Boldrey (1937) and Woolsey, Erickson, and Gilson (1979) used evoked potentials and electrical stimulation of the brain to find that human somatosensory and motor cortices vary by up to 100 percent.

Combinations of Brodmann's brain regions also show individual variation. Using morphometric and cytoarchitectonic techniques, Andrews, Halpern, and Purves (1997) found that the volume of the lateral geniculate nucleus, the cross-sectional area of the optic tract, and the surface area and volume of V1 all varied by about a factor of 2.5 across individuals.[23] The optic nerve, the optic tract, and areas V1, V2, and V3 reveal cortical magnification. That is, this portion of the visual system devotes disproportionately more cortex to the foveal regions of the retina than to the parafoveal regions. For areas V1, V2, and V3, as the distance from the fovea (eccentricity) increases, the surface area of the cortex dedicated to one degree of visual arc decreases. Dougherty et al. (2003) used fMRI data to document individual variation in the degree of cortical magnification. This suggests, as did the previous case, that properties of functionally significant combinations of brain regions, such as the "early visual system," might be multiply realized by properties/relations of different combinations of cells or perhaps by properties/relations of distinct combinations of cortical columns. Once again, though we have a case of type (iii), we nonetheless have strong evidence for individual variation at the level and the likelihood that the properties that are realizers will also be varied, hence making multiple realization likely.

Finally, there are apparent variations in whole brains. Individual variations in the weights, areas, and volumes of large brain structures have long been documented using postmortem analyses. More recently, however, MRI techniques have enabled investigators to measure individual variation in areas and volumes of brain regions in vivo, thereby avoiding confounding factors, such as variations due to differences in postmortem preservation techniques. Allen, Damasio, and Grabowski (2002) reported the results of MRI of brain volumes in normal subjects. All subjects, 23 male and 23 female, were right-handed with no left-handedness in first-degree relatives. There were healthy, with no history of neurological or psychiatric illness. They were primarily of European descent recruited primarily from the Iowa City community. The principal results regarding individual variation are shown in the table 22.3.

Table 22.3. Variations in Brain Dimensions in Humans

		Male		Female	
Region	Side	Mean cc (range)	s.d.	Mean cc. (range)	s.d.
Hemisphere (including cerebellum)					
	Left	618.9 (512.2-722.9)	54.7	547.0 (472.6-674.3)	47.8
	Right	621.7 (513.5-733.9)	58.3	552.7 (475.8-688.3)	49.5
Frontal					
	Left	205.2 (163.4-286.9)	25.8	182.8 (157.0-228.0)	17.0
	Right	208.3 (170.0-263.8)	24.0	186.2 (162.4-228.5)	17.2

The last four cases, at the very highest levels of organization, were of type (iii), where we lack detailed accounts of the specific lower level properties that realize properties at the particular higher level in question. However, as we have briefly noted in our discussion, it should now be clear that such examples still support the likelihood of multiple realization. However, to further illustrate how such examples support multiple realization, let us consider in more detail how the case at hand provides such evidence. Our examples of cortical columns, ocular dominance columns, Brodmann's brain regions and combinations of them, entire lobes, and entire brains have a common feature. The structures vary in size. This suggests that these structures differ in the number of neurons they contain.[24] Bare differences in numbers of cells might give rise to different combinations of properties/relations that realize properties at the level of cortical columns, ocular dominance columns, and so forth, and hence might well suffice for multiple realization. In addition, differences in the number of cells in a given region are also likely to induce or be correlated with new properties/relations among the cells. Differences in the numbers of cells will not necessarily be "just more of the same, only bigger." The latter point is rather speculative neuroscience which is why we have chosen to put the case into category (iii), wherein the properties of one level do not yet enable us to determine which properties are likely to be realizers of properties at the next higher neuroscientific level. However, once again, we have strong evidence for individual variation at this level and the levels below it, and thus the likelihood that the properties that are realizers will also be varied, hence making multiple realization likely.

Our brief sampling of empirical findings makes it plausible that there is a wide base of evidence supporting working neuroscientists in their twin commitments that there are a variety of neurobiological levels and that we find individual variation at many of these levels. Let us now briefly explain why these two commitments provide evidence for multiple realization, drawing together our foregoing points about the particular examples of types (i)–(iii).

As we have seen, depending on whether the examples are instances of cases (i), (ii), or (iii), the strength of evidence for multiple realization varies. Some psychological properties, such as being a deuteranope, are instances of situation (i) and sometimes (ii), and in these cases we have clear and compelling evidence for multiple realization. For other properties, we have found that our understanding is still in situations (ii) or (iii), and we are much less sure what the realization bases of these properties might be. Nonetheless, we have seen that even in cases of type (ii) and (iii) we still have evidence, albeit weaker in nature, for individual variation in the properties that are candidates for being realizers at the neurobiological levels below the relevant realized properties. Because the lower neurobiological levels displaying such individual variation supply the realizers for the psychological properties in question, and because we have found with our type (i) examples that in neurobiology individual variation favors multiple over univocal realization, such examples provide indirect support for the multiple realization of the relevant psychological properties. We can see that cases of type (i) strongly support the existence of multiple realization and that situations of type (ii), depending on their maturity, strongly or more weakly support multiple realization. Furthermore, we have now also seen that examples of type (iii) provide weaker support for multiple realization, within the context of our wider evidence that all higher level scientific properties are realized and our findings in cases of type (i) from neurobiology that individual variation at a level grounds multiple realization at that level.

Despite the differences in the strength of evidence the kinds of case supply, we have thus found that cases of types (i), (ii), and (iii) in neurobiology all provide support for the multiple realization of psychological properties at the relevant level. How one uses this evidence is a delicate issue. It should be obvious that potentially we have the grounds for a very blunt enumerative induction to the conclusion that *all* human psychological properties are multiple realized at *all* neurobiological levels. However, we should also be clear that such a blunt enumerative induction to such a strong conclusion is questionable, especially given the rough-and-ready nature of our sampling of empirical evidence and because as different areas become more mature we may well discover that *some* human psychological properties are univocally realized by the properties at *some* neurobiological levels. We therefore put this stronger conclusion to one side. A weaker conclusion, we contend, can be safely supported by such an induction and challenges the recent philosophical skepticism that there is *any* multiple realization in neurobiology. For given the evidence we have seen that cases of types (i)–(iii) provide for multiple realization we suggest an enumerative induction, or other argument, can be safely constructed to ground a very plausible case for MMR—that is, the claim that *many* human psychological properties are multiple realized at *many* neurobiological levels.

To summarize, we applied our more precise theoretical framework for scientific realization and multiple realization to a brief sampling of some of the evidence about individual variation in humans at the various levels of neuroscientific organization, from the biochemical to the entirety of the brain. We argued that our framework shows that the nature of our empirical evidence provides prima facie plausible reasons to believe that many human psychological properties are multiply realized at many neuroscientific levels. One might not think that such a conclusion is particular noteworthy or surprising given the nature of the empirical evidence so commonly found in neurobiology at all levels. But given the increasingly widespread contention of philosophers that there is *no* multiple realization in neurobiology, we hope our admittedly rough-and-ready defense of MMR is dialectically useful.

To put the point in Darwin's favored terminology: One can show, by a long catalog of facts, that parts whose properties must be called important for the realization of a psychological property often vary in the individuals of the same species. Consequently, we have evidence for multiple realization of many psychological properties, at many neurobiological levels, in different organisms of a single species. There presently look to be few levels of neuroscientific organization that are simple or uniform enough to provide for a univocal realization of many psychological properties. We thus have prima facie plausible empirical reasons to accept the truth of MMR.

3. MMR: Philosophical Objections and Wider Implications

As suggested in our introduction, the conclusions of the last section are not news to working scientists. For reasons that should now be clear, we contend that many neuroscientists find pervasive multiple realization to be a common and uncontentious feature of their subject matter (though obviously they do not use the phrase "multiple realization" to describe this phenomenon). However, because many philosophers deny the existence of multiple realization in neuroscience and argue that it is has radical or dangerous implications for scientific methodology, we want to conclude this chapter by considering some important philosophical objections to multiple realization. We hope to both further illuminate the nature and implications of MMR and also assess the substance of such concerns. We begin by considering older objections to multiple realization and work our way through more recent critiques.

3.1. Objection 1: The Evidence Used to Establish Multiple Realization Does Not Preclude Univocal Realization or Species-Specific Narrow Identities

This type of objection comes in two related flavors, based on a defense of either a species-specific univocal realization base for psychological properties or a so-called species-specific narrow identity between neural and psychological properties. Given the focus of the discussion on multiple realization, we structure our discussion around the former version of such an objection, though we finish by showing how our response also applies to the narrow identity version.

Crucially, such older objections were mounted in response to defenses of multiple realization, primarily some arguments of Hilary Putnam (1967), based around so-called philosophical thought experiments, such as imagining a silicon-based life form or robots, designed to show the possibility of multiple realization bases for psychological properties. However, in response, critics responded that they could just as easily imagine that psychological properties were uniquely realized by *some* lower level properties in a certain species.

This is apparently one of the points Jaegwon Kim is concerned to highlight when he observed that the existence of some diversity in lower level properties does not automatically preclude a univocal realization of higher level properties at some level:

> The fact that two brains are physico-chemically different does not entail that the two brains cannot be in the "same physico-chemical state." ... To argue that the human brain and the canine brain cannot be in the same brain state because of their different physico-chemical structure is like arguing that there can be no microphysical state underlying temperature because all kinds of objects with extremely diverse microphysical compositions can have the same temperature; or that water-solubility cannot have a microstructural "correlate" because both salt and sugar which differ a great deal from each other in atomic composition are soluble in water. If the human brain and the reptilian brain can be in the same "temperature state," why can they not be in the same "brain state," where this state is characterized in physico-chemical terms? (Kim, 1972, pp. 189–190)

Here we take two points to be at work. First, Kim presses the general point that if imagined possibilities guide us in such cases, the critics of multiple realization can also avail themselves of the possibilities we can conceive. Building on this general point, second, Kim presses that although we can imagine multiple realization, we can also further imagine that among this multiplicity at some neuroscientific level there is a *shared* neuroscientific base across these differences that allows for a univocal realization (and/or the base for a so-called narrow, species-specific identity) at this neurobiological level.

We agree with this much of Kim's view. Our earlier work highlighting the fact that all claims of multiple realization are indexed to particular realizer properties and levels supports Kim's conclusion. In earlier debates, defenders of multiple realization wished to use this phenomenon to undermine type-type identities between neural and psychological properties. However, to establish that there are no such identities, one needs to show that psychological properties are multiply realized by the properties *at all neuroscientific levels* and not just some. As Kim presses, although psychological properties might be multiply realized at neuroscientific level N the multiplicity of neuroscientific levels of organization still allows that at neuroscientific level (N – 1) there is a univocal realization for the diverse neuroscientific entities at level N and hence, by the transitivity of realization, for the relevant psychological properties. Exploiting this conclusion, Kim presses his first point about the methodology of imagined cases to suggest that because we can imagine such a situation, arguments from imagined cases of multiple realization actually fail to support the existence of the kind of multiple realization that would preclude any neural-psychological identities.

Looking more widely, as an overall strategy one can see the appeal of defending narrower claims about univocal realization or identities (see, e.g., Churchland, 1986, pp. 356–357). If we move from species-generic psychological properties to species-specific psychological properties, we are more likely, in a purely logical sense, to find univocal realizations. However, this logical point notwithstanding, we have now shown that our *empirical* evidence about individual variation within the human population plausibly undermines this route to univocal realization. Individual variations are variations among individuals of the same species, variations that are likely to be found even within much more narrowly circumscribed subpopulations of any species. The evidence for individual variations in neurophysiological and biochemical properties that we have used to support the MMR of psychological properties thus consequently provides prima facie plausible reasons to reject univocal realizations or species-specific narrow identities, even within organisms of a species.

The points that Kim makes about what we may or may not be able to imagine apparently does little to undermine the kind of argument from empirical evidence we have presented here. We happily admit that it is an open epistemic possibility that the MMR hypothesis may be undermined by future scientific discoveries, for like all claims supported on scientific grounds, we take this hypothesis to be empirically defeasible. In fairness, we should also emphasize that Kim's objection was originally directed against the very different, conceivability style defense of multiple realization presented by Putnam and not the type of argument we have built around actual empirical evidence. However, other defenses of univocal realization, both within organisms of certain species and across different species, have been defended by a variety of writers over the years (see, e.g., Bechtel and Mundale, 1999; Bickle, 2003; Braddon-Mitchell and Jackson, 1996; Churchland, 1986). At least some of these writers make claims about the actual world, and we suggest that our present evidence undermines their claims for species-specific (let alone transspecies) univocal realizations or narrow identities between human psychological and neural properties.[25]

3.2. Objection 2: The Dimensioned Accounts of Realization and Multiple Realization Trivialize the Issues

A rather different kind of objection comes from recent philosophical debates over the proper understanding of realization, and hence multiple realization, in the sciences. This objection claims that because the Dimensioned view of realization leads us to recognize the empirical support for MMR and hence endorse this thesis, this trivializes the whole question of multiple realization and undermines the Dimensioned view of realization itself. For these accounts obviously lead to accepting that there is far more multiple realization in the world than many philosophers of mind intuitively expected, and, the objection concludes, given these counterintuitive results we should therefore should abandon the Dimensioned view of realization and a key part of our argument for MMR.[26]

In response, it is worth marking that we have offered our accounts of scientific composition generally, and realization relations in particular, as parts of a wider understanding of the compositional concepts used in the sciences, in particular as views of the concepts deployed in mechanistic explanations in the special sciences. We suggest that the success of such accounts should consequently be judged by how well they do in capturing the features of the concepts actually used in such scientific explanations, and we have defended our theoretical framework with this point clearly in mind. What if the resulting account does not accord with some philosopher's intuitions or prior expectations about the nature or implications of compositional concepts? As in so many other cases where the sciences have surprised philosophers, we do not see why such a conflict would pose a problem for the success of our theoretical framework—for our work would simply highlight how our pretheoretic intuitions clash once again with the findings of the sciences and should be revised accordingly.

Furthermore, we should also emphasize that we take it to be a thoroughly empirical matter both whether we find multiple realization in some area of the sciences, and, if there is any, how much multiple realization there is. There may well be areas of the sciences where the properties of this area, Y_1–Y_n, are rarely if ever multiply realized by the properties, X_1–X_n, of the relevant lower level science. In such cases, the evidence, largely in the shape of the nature of the well-confirmed mechanistic explanations offered in these areas, will determine what one is justified in saying about these issues. What if such an examination of the relevant mechanistic explanations justified us in thinking that Y_1–Y_n are multiply realized, or even *pervasively* or *massively* multiply realized, by the properties X_1–X_n of the lower level science, and this was in conflict with some philosopher's prior expectations about the extent of multiple realization in this area of the sciences? Once again, we simply do not see why this would pose a problem, for it seems to be one more example in which the sciences have corrected a mistaken belief about the structure of the natural world.

Given these points, and our foregoing accounts of scientific realization, as well as our survey of the empirical evidence for individual variation at all neuroscientific levels, we suggest that the trivialization worry fails to pose a substantive difficulty for our argument for MMR.

3.3. Objection 3: The Present Immaturity of the Psychological Sciences Means That We Are Unable to Usefully Assess Whether Multiple Realization Exists

This objection begins from the present immaturity of many areas of the psychological sciences where it is still a legitimate question which kind of psychological entity figures in successful explanations. What higher level psychological properties are supposed to be the same in the face of all the lower level biological, chemical, and physical diversity? Are the properties supposed to include both cognitive and qualitative properties? Are they supposed to be those borrowed from folk psychology or those legitimated by scientific psychology? Are the psychological properties that form the basis of such a scientific psychology likely to be local, embedded properties or so-called extended properties? And so on. Given all these open questions, proceeds the objection, we cannot be sure whether psychological properties are indeed multiply realized, because we lack a well-confirmed account of the nature of such properties to begin with.

Our simplest response here is that almost however one answers these questions about the relevant psychological properties, one still arrives at the multiple realization of the posited psychological properties. If individual variation is a pervasive feature of nervous systems at virtually all levels, then many legitimate kinds of psychological property will turn out to be multiply realized at many neuroscientific levels. If, say, a putative property of suffering a particular kind of pain is realized in certain kinds of neuronal activities, and those neurons or their activities display individual variation in the ways in which they still each give rise to one and the same kind of pain, then that pain will be multiply realized. Similarly, mutatis mutandis for a putative property of believing that 2 + 2 = 4. The qualification "almost" is meant to mark that our claim is obviously empirically defeasible once more, thus there could be developments in the psychological sciences or neurosciences that would establish univocal realizations. However, our point is that given our present evidence such eventualities appear unlikely.[27]

3.4. Objection 4: The Worry about "Grains"—Arguments from Multiple Realization Only Succeed by Typing Neural and Psychological Entities Using Different Grains of Description

Bill Bechtel and Jennifer Mundale provide arguments against the existence of multiple realization, the most important of which we examine shortly. Having argued that psychological properties are not multiply realized, Bechtel and Mundale then ake on the further burden of explaining why so many writers have mistakenly hought they were. They tell us:

> one diagnosis of what has made the multiple realizability claim as plausible as it has been is that researchers have employed different grains of analysis in identifying psychological states and brain states, using a coarse grain to identify psychological states and a fine grain to differentiate brain states. Having invoked different grains, it is relatively easy to make a case for multiple realization. But if the grain size is kept constant, then the claim that psychological states are in fact multiply realized looks far less plausible. One can adopt either a coarse or a fine grain, but as long as one uses a comparable grain on both the brain and mind side, the mapping between them will be correspondingly systematic. (1999, p. 202)

he diagnosis Bechtel and Mundale offer is that recent defenders of multiple realation have been careless with their use of descriptions of the relevant cases from e neurosciences and hence been led astray.

We can begin to appreciate the basis of their diagnosis by noting that for any ate of affairs there are obviously a range of descriptions that all truly hold of it. For ample, we always have descriptions truly applying more general predicates as well other descriptions truly applying more specific predicates. Determinable and terminate predicates present one obvious case among many others: As we saw in tion 1, it is true, for instance, that both "the cone has the property of releasing a urotransmitter in response to light" and also that "the cone has the property of releasing a neurotransmitter in response to light in the neighborhood of 530 nm.” Because it will be important shortly, we also note that there is a very strong prima facie case, primarily based on considerations of ontological parsimony, that rather than these true sentences each picking out distinct properties, that is, the property of releasing a neurotransmitter in response to light *and* the property of releasing a neurotransmitter in response to light in the neighborhood of 530 nm, we should take both of these descriptions to be about the same property, in the property of releasing a neurotransmitter in response to light in the neighborhood of 530 nm. Some care needs to be used in discerning which of the true descriptions associated with our successful scientific explanations provide the most veracious guide to the relevant entities.

Bechtel and Mundale suggest that proponents of multiple realization may have been less than careful in addressing these kinds of issues in the scientific cases they use to support their claims. Defenders of multiple realization, these authors propose, may have gotten confused by using more specific descriptions with the lower level properties that are realizers but more general descriptions when describing the psychological properties they realize. The resulting diagnosis is thus that the appearance of homogeneity at the psychological level, in combination with heterogeneity at the level of the realizers, is merely an artifact of choosing different grains of description to apply to the properties at these levels, rather than an actual feature of the properties illuminated by our scientific explanations.

Against the background of Bechtel and Mundale's arguments against multiple realization, the grains diagnosis obviously has an important dialectical role. However, we must be careful to note that Bechtel and Mundale do not themselves use the grains point itself to argue against multiple realization, though such an argument is increasingly offered in conversation. The reader will already have appreciated that our response to the last criticism, underpinned by our earlier work, provides the basis of a blunt reaction to use of the grains point as an independent objection to multiple realization. Current neuroscience strongly suggests that at almost every level of organization the nervous systems of individuals in any significantly large population will display individual variation and hence we will have heterogeneity among the component entities that compose the same higher level entities. To see how the evidence we have laid out grounds such a response in detail, let us consider one body of findings outlined earlier, in the recent evidence about changes in dendritic spine structure within the same individual over time—though we should emphasize that similar problems arise with a wide range of different findings as we already illustrated in section 1 for the cone's properties.

In such cases, it is plausible that the relevant lower level realizer properties do indeed change, because these changes are tracked in the innovative experimental methodology. The proponent of such a grains critique might nonetheless respond that the relevant realizers across time are always instances of some common, general property, for example, the property of being a property of a dendritic spine. However, as the point about ontological parsimony makes plausible, this suggestion should be rejected on general grounds. For it would lead us to posit *two* sets of properties, the specific properties of dendritic spines outlined and also some further general property, such as the property of being a property of a dendritic spine or the like. There are obvious reasons from ontological parsimony to deny that there are such general properties, for instead we may simply assume there are two sets of specific and general predicates we may use to refer to the property referred to by the specific predicate.

Though we appear to have good reason to posit heterogeneity among the realizers, perhaps the proponent of such a grains critique may have better luck denying the needed homogeneity with the higher level property instances. Once again, however, our empirical evidence poses a problem for any such claim, because the cases at hand occur within the same individual who often appears to persist in having an instance of some specific psychological property. Thus unlike cases involving organisms of different species or victims of trauma, where one may raise real concerns about sameness of psychological property, there is far less scope for arguing that the psychological property must change as its realizers change. Absent some background theoretical reason for the assumption that realized property instances with differing realizers must be themselves instances of distinct properties, it thus appears that the nature of our empirical evidence in this kind of example supports the same psychological property being realized in the individual by different realizers over time.

Our brief examination of one body of empirical evidence highlights how the nature of such findings often frustrates the easy use of concerns about grains of descriptions as an independent argument against the existence of multiple realization. However, as we emphasized, Bechtel and Mundale only use the grains point *after* presenting their arguments against multiple realization, so let us finally turn to the most important of these criticisms.

3.5. Objection 5: The Existence of Multiple Realization Would Entail the Methodological Isolation of Psychology

This final objection has recently become quite influential (see Bechtel and Mundale, 1999, and Shapiro, 2004, among others). Such objections start by assuming that multiple realization entails that there is no intertheoretic constraint between the sciences studying realizer and realized properties. But it is then plausibly shown, using real examples, that neuroscience and psychology do intertheoretically constrain each other. Thus it is concluded that multiple realization does not exist for neural and psychological properties.

Elsewhere we dub this the methodological argument against multiple realization and critically examine it at length (Aizawa and Gillett, "Multiple realization and methodology in neuroscience and psychology," unpublished working paper). Of all the recent objections to multiple realization, we suspect this has been the most persuasive, not least because it links so seamlessly with the heated philosophical discussions of methodology and funding priorities for neuroscience and psychology. However, we believe that our earlier work quickly shows that the key premise of such objections is mistaken—it is simply false that multiple realization entails that there is no intertheoretic constraint between the sciences studying realizer and realized properties. In fact, the reverse is true, as our earlier remarks about the methodological implications of the transitivity of scientific realization suggested.

As we have seen from our examination of concrete scientific cases, certain properties are realizers of other properties *only if* the powers of these properties together noncausally determine the powers of the realized property instance. There are other important features of realization relations, but given only the latter aspect, we can see that there are very clear and tight constraints on which types of property are such that their instances can realize some other instance together. As a result of this ontological point, there are strong reasons to expect that under certain circumstances (such as having sufficiently well-confirmed theories), the nature of realization relations will ground intertheoretic constraints between the disciplines studying realizer and realized properties.

For example, if one has a very well-confirmed theory of the nature of some realized property, that is, an account of its individuating powers, then this theory can be used top-down to guide and even constrain research about the realizers of this property given other information about them. These realizers must result in the known powers of the realized property, so one can exclude certain hypotheses about the realizers or prioritize others, depending on whether these hypotheses make claims about the realizers' powers that together allow them to noncausally result in the powers of the realized property. In the reverse direction, working bottom-up, if one has a well-confirmed account of the nature of the realizer properties of some realized property, this constrains theories of the realized property in various ways. For instance, precise knowledge of the realizers' powers can exclude or prioritize certain hypotheses about the individuating powers of the realized property. Theories of the realized property's nature are in part plausible to the degree to which we can see that the powers the hypothesis accords to the realized property are such that they can noncausally result from the powers attributed to the realizers by our well-confirmed account of the latter.

Such conclusions are startling to many philosophers. Nonetheless, these abstract points ground a persuasive case that under certain conditions, because the properties involved in realization relations ontologically determine each other there will consequently be a range of obvious intertheoretic constraints between the disciplines studying realizer and realized properties. Because cases of multiple realization all involve realization relations, the obvious conclusion is that multiple realization, under the appropriate conditions (such as having suitably confirmed theories, etc.), actually results in the disciplines studying realizer and realized properties intertheoretically constraining each other. Given this point, rather than undermining the existence of multiple realization, all the evidence for intertheoretic constraint between neuroscience and psychology looks very different—such evidence is compatible with or even further supports the existence of multiple realization!

We can thus see that the methodological argument against multiple realization is ultimately unsound. Perhaps more important, we can take the sting out of recent concerns about the implications of multiple realization. For multiple realization, as a species of realization, simply does not entail, if it exists, that psychology will be methodologically isolated. In fact, the reverse is true. The existence of multiple realization increases the likelihood that one of the most fruitful approaches to research is a coevolutionary research strategy, where neurobiology and psychology each constrain each other in a mutually beneficial dance of fit and adjustment. We can therefore see that mistaken assumptions about the methodological implications of multiple realization have led to ungrounded fears about the impact of multiple realization and have apparently blinded many philosophers to what we have suggested is the wealth of evidence for MMR.

4. CONCLUSION

To finish, consider this passage, in which William Wimsatt bemoans the approach that has recently characterized specific areas of philosophy of science and wonders whether we can do better as philosophers. He states:

> We show our own disciplinary biases and force them on others: the various "philosophies of X" often seem to be more about arguments internal to philosophy than "of" anything....
>
> *Can we still be recognizably philosophical while letting the subjects of "philosophies of" shine through much more clearly and inspire new philosophies, rather than merely exporting the same old "philosophical" disputes to these new territories?* (Wimsatt, 2007, p.7; emphasis in original)

Though the focus of Wimsatt's remarks is slightly different than our own, we hope our work not only illuminates some of Wimsatt's concerns but also provides one positive answer to his important question.

We began the chapter by noting that researchers in the philosophy of psychology, and those in the philosophy of neuroscience, each have very different views about the existence of multiple realization and its methodological implications, as well as endorsing differing positions about connected issues such as the possibility of reduction. However, just as Wimsatt suggests, we have found that in this case the clashing positions in these two "philosophies of" appear to be philosophical artifacts, rather than being underpinned by the relevant areas of science. We have shown that neither the position in philosophy of neuroscience that there is no multiple realization nor the view in philosophy of psychology that multiple realization grounds the strong methodological autonomy of psychology is correct.

As we have seen, driven by the empirical findings they routinely confront, working neuroscientists endorse the existence of a range of neurobiological levels and the variation of individuals at all of these levels. By finally applying a more precise framework for realization and multiple realization to the empirical evidence that underpins the latter commitments of scientists, we have shown that many psychological properties of humans are plausibly multiply realized by properties at many neurobiological levels. In addition, however, we have also noted that such multiple realization produces intertheoretic constraint, rather than methodological autonomy, between the sciences that study realizer and realized properties—including psychology and the swath of disciplines that study neurobiological levels. Our work thus illuminates one important example where Wimsatt is correct in chiding researchers in "philosophies of," for we have shown that philosophers of neuroscience and philosophers of psychology have defended positions that fail to reflect either the nature of the disciplines they study or the commitments of the workers in these areas.

In addition, we hope that our substantive work here has been recognizably *philosophical* and provides one answer to Wimsatt's important question about how we may still pursue philosophy while being guided by scientific findings and commitments. Providing an abstract framework for scientific composition, as well as schemata for realization and multiple realization in the sciences, is a recognizably philosophical project—though clearly informed and guided by the sciences themselves—and one we have seen can produce substantive philosophical conclusions. We suggest that philosophers of science generally can profit by pursuing similar projects in what has come to be called the metaphysics of science: the careful, abstract investigation of ontological issues as they arise *within* the sciences and their findings, models, explanations, and so on. Our conclusion is therefore that even philosophers of science engaged in "philosophies of" this or that science may benefit by adding the tools of the metaphysics of science to their methodological armory.

ACKNOWLEDGMENTS

Thanks to Marica Bernstein for comments on an earlier draft of this chapter.

NOTES

1. There are, of course, a number of stronger but related versions of this hypothesis. For example, one can change the two "many" quantifiers to "most" or "all" and not limit the hypothesis merely to human psychological processes. For the space of the present discussion, however, we defend only this weaker hypothesis. Furthermore, as we see later, properties usually come along in packages with distinctive powers, individuals, and processes. Consequently, as well as the multiple realization of properties, we also endorse the multiple constitution and multiple implementation of psychological individuals and processes in humans. Once again, however, we leave that claim to one side here. Finally, we also note that we defend multiple realization, constitution, and implementation for the psychological properties, individuals, and processes found in most terrestrial species, but again we limit our focus to humans here.

2. Exactly which researchers hold which views, and hence narratives, is a very contentious issue in this area, for reasons that will become apparent. However, though we accept that some philosophers of psychology have defended strong methodological autonomy for psychology, and hence endorse this narrative, we argue in Aizawa and Gillett ("Multiple realization and methodology in neuroscience and psychology," unpublished working paper) that this was plausibly not the position of Fodor (1968), or other writers defending the received view of special sciences such as William Wimsatt or Philip Kitcher.

3. Historically speaking, philosophers of mind have more often discussed multiple realizability, rather than multiple realization. Here we limit our attention to multiple realization, because it allows us to sidestep issues about the proper modality, it simplifies and abbreviates further discussion, and once multiple realization is established, the multiple realizability follows rather easily.

4. For a more extensive exposition of the nature of these compositional concepts, see Gillett ("Making sense of levels in the sciences: Composing powers, properties, parts, and processes," unpublished working paper).

5. It is worthwhile marking some of the distinct notions of realization current in recent philosophical discussions and that the parties to the present debate only defend variants of one of these concepts. (See Endicott, 2005, for a survey of some of main varieties of the concept of realization.) Some writers fail to appreciate the variety of concepts of realization, and hence confuse them, so it is worth distinguishing three kinds of notion of realization.

First, there is a group of semantic notions that we term *Linguistic*, or L-realization and that hold between entities in the world and some set of sentences. Famously, for example, the work of David Lewis on topic-neutral Ramseyfication and theoretical terms uses a notion of L-realization. Basically, on Lewis's view of realization it holds between entities in the world and the set of Ramseyfied sentences putatively defining some theoretical term F—crudely put, an entity X L-realizes F when the entity X satisfies the relevant Ramsey sentences for F.

Second, there is a kind of computational or mathematical relation commonly referred to as realization and used in both the sciences and philosophy, which we call *Abstract* or A-realization. Very crudely, X is taken to A-realize Y if the elements of X map onto or are isomorphic with the elements of Y. This notion of realization is commonly utilized with formal models and hence with work using such models, for example in computational accounts of cognitive processes. Note that here the relata of such realization relations are largely unconstrained because A-realization simply holds in virtue of a mathematical mapping, or isomorphism, which can obviously hold between all manner of entities.

Finally, the third kind of notion of realization is what we may call *Causal-Mechanist* or M-realization. The latter contrasts with L- and A-realization by having as relata causally individuated entities in the world, often (though not exclusively) property instances. M-realization has been the focus of many writers, and in particular philosophers of science have been especially interested in such relations, which they take to be posited in so-called mechanistic explanations in a range of the special sciences. In our discussion, for obvious reasons, we focus exclusively on a notion of M-realization when we discuss realization.

6. There are three types of cone photopigments. Sometimes they are classified as blue, green, and red and sometimes they are classified as short wavelength, middle wavelength, or long wavelength (S, M, and L), respectively.

7. In actual neuroscience, it is common to decompose a cone photopigment molecule into a protein component and a nonprotein chromophore. The protein component might then be decomposed into distinct amino acids. We skip these intermediate levels between

the entire cone photopigment molecule and the individual atoms primarily because of the greater familiarity of atomic-molecular relations and to simplify the exposition.

8. This thumbnail account is defended in Gillett (2002, 2003). A full account of the Dimensioned view of realization, as a part of integrated view of the compositional relations posited in the sciences between packages of powers, properties, individuals and mechanisms, is offered in Gillett ("Making sense of levels in the sciences," unpublished working paper).

9. Scientific purists will no doubt observe the simplification we have been working with in referring to the property of being maximally sensitive light of 530 nm. Distinct experimental methods yield slightly different values for maximal sensitivity. Moreover, these methods include one or another measure of error. What remains through this simplification is that current science takes the maximum absorption spectra of distinct amino acid sequences to be the same up to experimental error.

10. To avoid confusion, we reiterate what might seem obvious: Our theory schema is intended to follow the usual convention and illustrate how properties are multiply realized—for if two instances of the same property are differently realized at the same level, then the property is *multiply* realized. However, as we see in examples to follow, an instance of a higher level property may itself also be multiply realized over time by distinct lower level realizers. In such cases we thus have both a higher level property and also a single instance of this property that are multiply realized. However, for simplicity, we focus primarily on properties as the entities that are multiply realized.

11. Some readers may be concerned that adding (iv) leaves us with a dangerously imprecise notion in that of a level of entities. However, as William Wimsatt and others have argued (Wimsatt, 1976, 1994), there is a reasonably clear scientific notion of a level of entities, under some condition, as entities that do or can participate in the same causal mechanisms under those conditions (or which participate in processes that together implement other processes). This scientific concept of a level can underwrite (iv), and elsewhere one of us has also outlined a precise definition of this notion of a level (Gillett, "Making sense of levels in the sciences," unpublished working paper).

12. As we mentioned, distinct sets of atoms with distinct properties can give rise to distinct molecules with instances of the same property of absorbing a certain spectrum.

13. The methodological implications of realization, and other compositional relations in the sciences, are also explored in more detail in Aizawa and Gillett ("Multiple realization and methodology in neuroscience and psychology," unpublished working paper).

14. See Sharpe et al. (1990).

15. The astute reader will no doubt again anticipate the concerns that arise about these examples with regard to which grains of description, specific or general, one should take to pick out the actual properties. We happily agree that these are substantive issues, and interesting cases can be presented for a number of interpretations. However, given the range of lower level individual variation, in the same manner that we highlighted in the molecular-to-cellular example, we suggest that it is highly plausible that in the final analysis at least *some* of these findings will again ultimately support multiple realization in the cell-to-tissue example.

16. This is one of the central contentions of Aizawa (2007).

17. As suggested previously in note 7, in actual practice it is common enough for neuroscientists to relate amino acid sequences and their properties/relations and powers, on the one hand, to proteins and their properties/relations and powers, on the other, rather than relating atoms, atomic-level properties/relations, and atomic-level powers to proteins, protein-level properties/relations, and protein-level powers. However, this practice changes nothing philosophically relevant to our concerns.

18. For simplicity, we set aside properties that other components of a dendritic spine contribute to the shape of the spine.

19. Of course, we are setting aside mention of the contributions of other properties of the subcellular components of the neuron. This is merely for expository simplicity.

20. See, for example, Zuo, Lin, Chang, and Gan (2005) and Majewska, Newton, and Sur (2006).

21. Perhaps glial cells should be included as part of the realizer of a cortical column, but merely for the sake of simplicity of exposition we set the consideration aside. Nothing of philosophical import appears to turn on it.

22. Van Essen, Newsome, and Maunsell (1984) found almost as much variation in the surface area of V1 in the macaque. These variations appear to be independent of body weight. In a study with only six macaques, Horton and Hocking (1996) found much less variability.

23. Of course, the optic tract and lateral geniculate nucleus are not cortical structures.

24. Here we again set aside complications regarding the possible role of glial cells.

25. Aizawa (2007) further defends the claim that essentially all psychological properties are multiply realized in virtue of the variations in amino acid sequences that are components of all nerve cells, and indeed all major tissues of the body. That paper presupposes something like the Dimensioned view of realization and theory of multiple realization advanced here.

26. A number of philosophers have pressed concerns of this kind in conversation and talks.

27. As a side note, it is worth examining a related worry analytic philosophers often raise: Why doesn't the logical possibility of individuating psychological properties just as finely as their putative realizers immediately falsify our claim? Once again, we note that we take successful scientific explanations to be our guide to the entities that exist, whether powers, properties, individuals, or processes, and whether such entities should be taken to be psychological. Thus, although individuating psychological properties so finely might be thought to be possible, this would still be a long way from showing properties individuated in such a manner underlie successful scientific explanations. Thus, by itself, the possibility does not pose a problem for our claim.

REFERENCES

Aizawa, K. (2007). The biochemistry of memory consolidation: A model system for the philosophy of mind. *Synthese* 155: 65–98.

Aizawa, K., and Gillett, C. (unpublished working paper). Multiple realization and methodology in neuroscience and psychology.

Allen, J. S., Damasio, H., and Grabowski, T. J. (2002). Normal neuroanatomical variation in the human brain: An MRI-volumetric study. *American Journal of Physical Anthropology* 118: 341–358.

Andrews, T., Halpern, S., and Purves, D. (1997). Correlated size variations in human visual cortex, lateral geniculate nucleus, and optic tract. *Journal of Neuroscience* 17: 2859–2868.

Bechtel, W., and Mundale, J. (1999). Revisiting multiple realization. *Philosophy of Science* 66: 175–205.

Bickle, J. (2003). *Philosophy and Neuroscience: A Ruthlessly Reductive Account*. Boston: Kluwer.

Blake, R., and Sekuler, R. (2002). *Perception*, 4th ed. Boston: McGraw-Hill.[a]

Braddon-Mitchell, D., and Jackson, F. (1996). *Philosophy of Mind and Cognition: An Introduction*. Oxford: Blackwell.

Brodmann, K. (1909). *Vergleichende Lokalisationslehre der Grosshirnrinde in ihren Principien, dargestellt auf grund des Zellenbaues*. Leipzig: Johann Ambrosius Barth Verlag. 2nd ed., 1925. English translation by Laurence J. Garey, *Brodmann's 'Localisation in the Cerebral Cortex'* (Smith-Gordon, 1994); new impression, Imperial College Press, 1999.

Churchland, P. S. (1986) *Neurophilosophy: Towards a Unified Science of the Mind/Brain*. Cambridge, Mass.: MIT Press.

Darwin, C. (1964). *On the Origin of Species. A Facsimile of the first edition*. Cambridge, Mass.: Harvard University Press.

Dougherty, R., Koch, V., Brewer, A., Fischer, B., Modersitzki, J., and Wandell, B. (2003). Visual field representations and locations of visual areas V1/2/3 in human visual cortex. *Journal of Vision* 3: 586–598.

Endicott, R. (2005). Multiple realizability. In *The Encyclopedia of Philosophy*, 2nd ed. New York: Macmillan.

Fodor, J. (1968). *Psychological Explanation*. New York: Random House.

Gillett, C. (2002). The dimensions of realization: A critique of the standard view. *Analysis* 62: 316–323.

Gillett, C. (2003). The metaphysics of realization, multiple realizability and the special sciences. *Journal of Philosophy* 100: 591–603.

Horton, J., and Hocking, D. (1996). Intrinsic variability of ocular dominance column periodicity in normal macaque monkeys. *Journal of Neuroscience* 15: 7228–7339.

Kaschube, M., Wolf, F., Geisel, T., and Löwel, S. (2002). Genetic influence on quantitative features of neocortical architecture. *Journal of Neuroscience* 22: 7206–7217.

Kaschube, M., Wolf, F., Puhlman, M., Rathjen, S., Schmidt, K., Geisel, T., et al. (2003). The pattern of ocular dominance columns in cat primary visual cortex: Intra- and interindividual variability of column spacing and its dependence on genetic background. *European Journal of Neuroscience* 18: 3251–3266.

Kim, J. (1972). Phenomenal properties, psychological laws, and identity theory. *Monist* 56: 177–192.

Majewska, A. K., Newton, J. R., and Sur, M. (2006). Remodeling of synaptic structure in sensory cortical areas in vivo. *Journal of Neuroscience* 26: 3021–3029.

McKinney, R. A. (2005). Physiological roles of spine motility: Development, plasticity, and disorders. *Biochemical Society Transactions* 33: 1299–1302.

Merbs, S., and Nathans, J. (1993). Absorption spectra of the hybrid pigments responsible for anomalous color vision. *Science* 258: 464–466.[b]

Nagel, E. (1961). *The Structure of Science*. New York: Harcourt, Brace, and World.

Nathans, J., Thomas, D., and Hogness, D. S. (1986). Molecular genetics of human color vision: The genes encoding blue, green, and red pigments. *Science* 232: 193–202.

Neitz, M., and Neitz, J. (1995). Numbers and ratios of visual pigment genes for normal red-green color vision. *Science* 267: 1013–1016.

Penfield, W., and Boldrey, E. (1937). Somatic, motor and sensory representation in the cerebral cortex of man as studied by electrical stimulation. *Brain* 60: 389–443.

Peters, M., Jänke, L., Staiger, J. F., Schlaug, G., Huang, Y., and Steinmetz, H. (1998). Unsolved problems in comparing brain sizes in *Homo sapiens*. *Brain and Cognition* 37: 254–285.

Putnam, H. (1967). Psychological predicates. In W. H. Capitan and D. D. Merrill (Eds.), *Art, Mind, and Religion*. Pittsburgh: University of Pittsburgh Press, 37–48.

Rademacher, J., Morosan, P., Schormann, T., Schleicher, A., Werner, C., Freund, H.-J., et al. (2001). Probabilistic mapping and volume measurement of human primary auditory cortex. *NeuroImage* 13: 669–683.

Shapiro, L. (2004). *The Mind Incarnate*. Cambridge, Mass.: MIT Press.

Sharpe, L. T., Stockman, A., Jägle, H., and Nathans, J. (1999). Opsin genes, cone photopigments, color vision, and color blindness. In K. R. Gegenfurtner and L. T. Sharpe (Eds.), *Color Vision: From Genes to Perception*. Cambridge: Cambridge University Press.

Shoemaker, S. (1980). Causality and properties. In P. van Inwagen (Ed.), *Time and Cause*. Dordrecht: Reidel.

Sorra, K., and Harris, K. (2000). Overview on the structure, composition, function, development, and plasticity of hippocampal dendritic spines. *Hippocampus* 10: 501–511.[c]

Stensaas, S. S., Eddington, D. K., and Dobelle, W. H. (1974). The topography and variability of the primary visual cortex in man. *Journal of Neurosurgery* 40: 747–755.

Van Essen, D., Newsome, W., and Maunsell, J. (1984). The visual field representation in striate cortex of the macaque monkey: Asymmetries, anisotropies and individual variability. *Vision Research* 24: 429–448.

Wässle, H. (2004). Parallel processing in the mammalian retina. *Nature Reviews Neuroscience* 5: 1–11.[d]

Wimsatt, W. (1976). Reductionism, levels of organization and the mind-body problem. In G. Globus, I. Savodnik, and G. Maxwell (Eds.), *Consciousness and the Brain*. New York: Plenum, 199–267.

Wimsatt, W. (1994). The ontology of complex systems: Levels of organization, perspectives and causal thickets. *Canadian Journal of Philosophy*, supplement 20: 207–274.

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

Woolsey, C., Erickson, T., and Gilson, W. (1979). Localization in somatic sensory and motor areas of human cerebral cortex as determined by direct recording of evoked potentials and electrical stimulation. *Journal of Neurosurgery* 51: 476–506.

Zuo, Y., Lin, A., Chang, P., and Gan, W. (2005). Development of long-term dendritic spine stability in diverse regions of cerebral cortex. *Neuron* 46: 181–189.