Abstraction and Explanatory Relevance, or Why Do the Special Sciences Exist?

Matthew C. Haug
The College of William & Mary

Abstract

Non-reductive physicalists have long held that the special sciences offer explanations of some phenomena that are objectively superior to physical explanations. This explanatory "autonomy" has largely been based on the multiple realizability argument. Recently, in the face of the local reduction and disjunctive property responses to multiple realizability, some defenders of non-reductive physicalism have suggested that autonomy can be grounded merely in human cognitive limitations. In this paper, I argue that this is mistaken. By distinguishing between two kinds of abstraction I show that the greater explanatory *relevance* of some special science predicates (to certain explananda) is both non-anthropocentric and not solely based on considerations of multiple realizability. This shows that the explanatory autonomy of the special sciences is safe from the local reduction and disjunctive property strategies, even if they are successful responses to the multiple realizability argument.

1. Introduction

For roughly forty years, non-reductive physicalism (NRP) has been the dominant view in philosophy of mind, and multiple realizability has been the dominant argument for it. Of course, there have long been prominent dissenters from this consensus who instead defend a version of reductive physicalism (RP), often by either (i) appealing to disjunctive properties or (ii) invoking the "local reduction" strategy, each of which I discuss below. I take the following three claims to be endorsed by a strong version of RP:

(R1) Biological, psychological or other special science properties are type-identical to physical properties.[1]

(R2) These property identities, together with the generalizations of physical theory, can offer explanations of why the generalizations of the special sciences are true.

(R3) Explanations that use special science predicates can be replaced by explanations that use only physical predicates without any loss of explanatory power.

I take it that the claim that the special sciences are "autonomous" is largely captured by the denial of (R2) and (R3). Strong NRP denies each of (R1)-(R3). However, in the last decade, a weaker version

---

[1] In this paper, I use the term 'physical' in a restricted sense to refer to low-level or microphysical properties and predicates which exclude mental, social, and some biological properties and predicates but include some neurochemical ones.

of (what is claimed to be) NRP has been suggested—call it "weak NRP"—which accepts (R1) (or at least does not argue against it) but denies (R2) and (R3) (at least on one reading of these claims) (see, e.g., Antony 1999, 18; Clapp 2001). Weak NRP construes reducibility as entirely, or primarily, an epistemic-cum-semantic issue.

In this paper, I will largely set aside questions about property reduction. That is, I will grant that the multiple realizability argument does not show that (R1) is false—that it does not establish strong NRP. Instead I will focus on the nature and status of weak NRP. If (R2) and (R3) are false, is this merely because of our cognitive limitations? Or are there non-anthropocentric reasons why the special sciences exist? After discussing how the local reduction and disjunctive strategies attempt to defend (R1) against the multiple realizability argument (Section 2), I argue that weak NRP is distinct from RP only if (R2) and (R3) are false for non-anthropocentric reasons (Section 3). I then identify such a reason—the greater explanatory *relevance* of special science predicates to some explananda. In Sections 4 and 5, I argue that distinguishing between two kinds of *abstraction* shows that this explanatory virtue is both non-anthropocentric and independent of considerations of multiple realizability, which blocks the local reduction and disjunctive strategies.

As is standard in debates about reductionism and mental causation, I assume a "sparse" conception of properties according to which differences in properties must reflect differences in causal powers (cf., e.g., Kim 1998, 105-6). Properties are individuated more coarsely on sparse views than on an extreme "abundant" conception according to which only synonymous or logically equivalent predicates express the same property. Because of this, a single sparse property may be *multiply expressible*: it may be expressed by several non-synonymous predicates. A certain kind of multiple expressibility will be important in Section 5.

2. The Multiple Realizability Argument and Two Responses

Here is one formulation of the multiple realizability argument against (R1), as applied to mental properties:

(1) If mental property *M* is reducible, then there is some natural, physical property *P* such that necessarily, for all *x*, *Mx* iff *Px*. ("General" or "uniform" property reduction is a necessary condition for reducibility.)

(2) For any natural, physical property *P* that realizes *M*, it is possible for some individual to lack *P* but have *M*, since some other natural, physical property, *Q*, could realize *M*.[2] (*M* is multiply realizable.)

(3) The disjunctive property, $P \vee Q$, is not a natural physical property.

(4) So, there is no natural physical property *P* such that necessarily, *Mx* iff *Px*.

(5) So, *M* is irreducible.

Many philosophers have claimed that this argument for property irreducibility fails. The most influential of these challenges—the local reduction strategy and the disjunctive strategy—deny premises (1) and (3), respectively.[3]

2.1. The Local Reduction Strategy

Put in schematic form, the local reduction strategy claims that reduction does not require that necessarily, for all *x*, *Mx* iff *Px*; rather, all it requires is that, given existing species or structure-types $S_1, S_2, ...$ necessarily, for all *x* that are S_1, *Mx* iff P_1x, and necessarily, for all *x* that are S_2, *Mx* iff P_2x, etc.

The most influential version of this strategy, what I call the eliminative version, concedes that *general* mental properties (that could be said to apply to many different species of organism) may be multiply realizable and hence irreducible. However, such a general, uniform reduction is not needed, according to this strategy, because these general mental properties are not natural. Rather, *being in pain* is merely a property in the "abundant sense"—is merely a concept or property designator—and is not projectible or causally efficacious. Further, the local reductionist argues that, perhaps due to

---

[2] Note that we are interested in *total* realizers since they are the only plausible candidates for the reduction of special science properties.

[3] Premise (2) has also been denied, by what I call the "direct argument" against the (substantive) multiple realizability of special science properties (cf. Shapiro 2004). I believe that what I say in §§4 and 5 could also be used to reply to this strategy.

evolutionary processes, the predicate 'is in pain' likely expresses a unique physical realizer in all normal members of a given species. Thus, the allegedly unnatural property *being in pain* is "partitioned" into natural properties that are unique to particular species or structural types: *pain-in-S_1*, *pain-in-S_2*, etc. Then, each of these species-specific high-level properties can be reduced to the (allegedly) unique physical property, P_1, P_2, ... that realizes it in the respective species. This version captures the local reduction strategy as it is most often presented by David Lewis (e.g. 1994, 304-308) and Jaegwon Kim (e.g., 1998, 109-111).[4]

2.2. The Disjunctive Strategy

The disjunctive strategy challenges premise (3) in the multiple realizability argument. In effect, it claims that proponents of the multiple realizability argument have simply focused on the wrong physical properties. According to the disjunctive strategy, there is a physical property that realizes M (namely, the disjunction of all of the possible non-disjunctive total realizers of M, which I denote by 'vP_i') for which it holds that necessarily, for all x, Mx iff vP_ix. With the threat of multiple realizability neutralized, the idea is that this lawlike correlation can be easily "enhanced" into a property identity (see Kim 1998, 97).[5]

3. Problems with an Anthropocentric Defense of Weak NRP

For years, many defenders of NRP followed Hilary Putnam and Jerry Fodor in thinking that the disjunctive strategy could be dismissed as a non-starter since all merely disjunctive properties (like *being a raven or a writing desk*) were supposedly "wildly heterogeneous," non-natural and (hence) unprojectible. However, as Kim (1998, 1999) has argued, this claim poses a problem for

---

[4] An alternative, *conservative* version of the local reduction strategy takes it to be part of a package with the disjunctive strategy. Each of the species-specific high-level properties is still identified with the corresponding physical realizer, but species-independent high-level properties are identified with disjunctive physical properties.

[5] See also Sober (1999). Although the disjunctive strategy is sometimes attributed to Kim (e.g. Antony 1999, 3-4), it is not clear that he ever *fully* endorsed it as a response to the multiple realizability argument. He seems to be ambivalent about the existence of disjunctive properties, and his position has changed over time.

NRP. Given that any multiply realizable property is necessarily co-extensive with the disjunction of its possible realizers, *if* the disjunctive property is unnatural and unprojectible, *then so is the multiply realizable property* (cf. Antony 1999, 13).

In response to this problem, there is a growing trend (even among defenders of NRP) of admitting that some disjunctive properties are natural (see Clapp 2001, Shoemaker 2007). If one takes this line, what is the upshot for NRP? In order to maintain the autonomy of psychology it appeared that we needed to show that mental properties were distinct from any physical properties that realize them. Now, in accepting the cogency of the disjunctive strategy, we have admitted that multiple realizable properties can be identified with the disjunction of their possible realizers (i.e. accepted (R1)). Does this mean that psychology is no longer an autonomous science—that we should also accept (R2) and (R3)?

Lenny Clapp has in effect suggested that this depends on how we interpret the word 'can' in (R2) and (R3). In response to the disjunctive strategy, Clapp suggests a

> … weaker version of NRP [that] does not deny that it is in principle possible for "ideal" scientists to formulate physicalistic predicates that would reduce our mentalistic predicates. … It rather claims that we really shall not and cannot [in practice] reduce our mentalistic predicates to physicalistic predicates. (2001, 135)

Clapp claims that weak NRP "suggests that the nonreducibility of mentalistic predicates is purely due to *our own* epistemological limitations" (ibid., italics added); we can formulate neither infinite disjunctive predicates nor the complicated "wide" individual physical disjuncts that the local reduction strategy will arguably need to invoke (ibid., 134; see also Antony 1999, 15).

However, denying (R2) and (R3) for purely *anthropocentric* epistemological reasons, as Clapp suggests, trivializes the debate between NRP and RP. On Clapp's suggestion, weak readings of (R2) and (R3) are false *not* because of the structure of the world but only because of our cognitive limitations. However, this seems to be something that even some prominent defenders of RP accept. For instance, in a series of papers and books, Jaegwon Kim (1998, 1999) defends the local reduction

strategy: there are no special science properties distinct from physical properties; (R1) is true. However, he of course admits that mental concepts or predicates are distinct from physical ones, *and* he admits that these mental concepts "may play a practically indispensable role in *our* discourse, both ordinary and scientific" (1999, 17, italics added; see also Kim 1998, 104-5).

These quotations suggest that what undermines the distinction between RP and weak NRP is not that the latter is defended on epistemic grounds but that it is defended on merely *anthropocentric* epistemic grounds. Thus, if we can identify some non-anthropocentric explanatory reasons why (R2) and (R3) are false, then, even though weak NRP remains primarily an epistemic doctrine, it will not collapse into RP. For, the defender of weak NRP can now assert (while the defender of RP will deny) that the reductive explanations promised by (R2) and (R3) are impossible *in principle* (i.e., even for epistemically ideal creatures) and not merely in practice.

4. The Generality and Explanatory Relevance of Special Science Predicates

There are at least two non-anthropocentric reasons to prefer explanations that use a special science predicate, *M*, to those that use physical predicates that express the individual physical realizers of *M*.

First, recall that Hilary Putnam's and Jerry Fodor's presentations of the multiple realizability argument emphasized the greater *generality* of special science predicates as compared to physical predicates. For example, 'is money' is satisfied by a wider variety of objects than any complicated physical predicate that describes a particular kind of money (see Fodor 1974). This greater generality is clearly an explanatory virtue that is not dependent on human cognitive limitations. To take another example: even if we were omniscient creatures, it would still be true that the predicate

'is a rigid, 1" square peg' applies to a greater variety of objects than the physical predicate that expresses one of its total physical realizers (see Putnam 1975).[6]

Second, defenders of the autonomy of the special sciences have often emphasized that special science predicates single out factors that are causally or explanatorily *relevant* to a given phenomenon, while derivations or generalizations that use physical predicates to try to explain that phenomenon will include many irrelevant details and hence fail to be explanatory. Call this the "relevance advantage." For example, Philip Kitcher claims that a derivation of the general principles of classical genetics from molecular genetics is not explanatory because "in charting the details of the molecular rearrangements the derivation would only blur the outline of a simple cytological story, adding a welter of irrelevant detail" (1984, 347).

Unfortunately, greater *generality* is not an explanatory virtue that a special science predicate, *M*, has over a disjunctive predicate that includes disjuncts for every possible realizer of *M*. So, once we have given up the multiple realizability argument as a defense of strong NRP, we can no longer appeal to greater generality in an argument against (R2) and (R3). What about the relevance advantage? Does it survive the disjunctive and local reduction strategies?

It is tempting to think that the greater relevance of special predicates is also compromised once we have given up the multiple realizability argument. Multiply realizable predicates abstract away from specific details about physical implementation, and, it might be thought, it is *because* of their greater abstraction that special science predicates are *both* more general than particular physical predicates *and* capture relevant "high-level" uniformities that are missed at the physical level. After all, it is plausible that if predicate P is more *abstract* than predicate Q, then P is more *general* than Q. So, if being *more abstract* is a necessary condition for the relevance advantage, then being *more*

---

[6] Some might claim that I am being somewhat sloppy here (and below) in writing of realizers of special science *predicates* instead of special-science *properties*. However, once we adopt weak NRP and accept (R1), realization arguably becomes a semantic relation.

general is a necessary condition for it as well. If this line of thought were correct, then, since the disjunctive and local reduction strategies undermine the *greater generality* of (natural) special science predicates, they would also undermine defenses of (R2) and (R3) that are based on the explanatory *relevance* of special science predicates. Given the prima facie plausibility of this line of thought, we should expect to see defenses of the relevance advantage trying to establish the multiple realizability (and hence generality) of special science predicates.

This is exactly the way that Kitcher defends the relevance of cytological predicates (and the irrelevance of molecular predicates) to the explanation of the transmission law of classical genetics. He claims in effect that (R2) is false because "explaining the transmission law requires identifying PS [pair separation]-processes as forming a natural kind to which processes of meiosis belong, and … PS-processes cannot be identified as a kind from the molecular point of view" (Kitcher 1984, 349). This is because PS-processes are "heterogeneous from the molecular point of view" (ibid.); they are "realized in a motley of molecular ways" (ibid., 350). Importantly, not only does Kitcher use multiple realizability/generality to defend the explanatory relevance of special science predicates, he also uses it to defend the *objective*, non-anthropocentric superiority of high-level explanations (ibid.).

I believe it is mistaken to tie the explanatory relevance of special science predicates to their multiple realizability/generality. The above argument that uses abstraction as a link to attempt to show that generality is a necessary condition for the relevance advantage is flawed. It conflates two kinds of abstraction: what I call *homotopic* and *heterotopic* abstraction. Only homotopic abstraction requires greater generality (or multiple realizability), and it is *not* required for the relevance advantage.

Predicates for determinable properties and those for their various determinates provide a paradigm case of *homotopic abstraction*. Determinates of a single determinable differ with respect to *aspects* of the determinable itself. For example, one plausible account of perceived colors sees them

as individuated by three aspects: hue, saturation, and brightness. A determinate shade of *red*, like

scarlet, is concrete in that it has a *particular* value of hue, saturation, and brightness. The

determinable *red* is abstract in that it is less specific regarding these aspects; it is characterized by a

range of hue, saturation, and brightness values. We can say that predicate X is *homotopically*

abstract relative to predicate Y when X takes up a larger volume of the same aspect space common

to both X and Y. Clearly, if predicate P is more *homotopically* abstract than predicate Q, then P is

more general than Q. (In addition to colors, think of the predicates 'has mass', 'has mass between 1

and 2 kg', and 'has mass of exactly 1.5 kg'.)

By contrast, the relation between a special science predicate and a physical predicate that

expresses one of its physical realizers differs from the relation between determinables and

determinates; special science predicates and physical predicates belong to *different* aspect spaces.[7]

For example, 'is red' will be exhaustively characterized by hue, saturation, and brightness values, but

a predicate for one of its physical realizers—say, one expressing a micro-based property of a robin's

breast[8]—will also be characterized by physical aspects like charge, diffraction coefficients, etc. Let's

say that X is *heterotopically abstract* relative to Y when the aspect space of X has fewer dimensions

than Y's aspect space, that is, when X is characterized by fewer aspects than Y.

With this distinction in hand, we can now show that generality is not necessary for the

relevance advantage. This is because (a) predicate *P* can be more heterotopically abstract than

predicate *Q without* being more general than *Q*, and (b) homotopic abstraction is not necessary for

the relevance advantage; heterotopic abstraction grounds it as well.

---

[7] For a related point, see Haug 2010. For the sake of argument, I assume that the aspects that
characterize physical predicates *include* those that characterize special science predicates, not that
special science and physical predicates belong to *disjoint* aspect spaces.
[8] A micro-based property is the property of having proper parts that are propertied and related in
certain ways.

Regarding (a), even if 'is in pain' were only realizable by a single physical property, say, *having C-fibers firing in a nervous system of an appropriate kind* (or a disjunctive property with this as one of its disjuncts), it would still be the case that 'is in pain' is characterized by fewer aspects— e.g., intensity, duration, and a variety of affective-motivational and sensory-discriminative aspects— than the physical predicate, since the latter will also be characterized by aspects like conductivity and degree of myelination.

Support for (b) – that homotopic abstraction is not necessary for the relevance advantage – can be extracted from an interesting footnote from Putnam 1975.

> Even if it were not physically possible to realize human psychology in a creature made of anything but the usual protoplasm, DNA, etc., it would still not be correct to say that psychological states are identical with their physical realizations. For, as will be argued below, such an identification has no *explanatory* value *in psychology*. (1975, 293, italics in original)

This quotation *directly* denies that greater *generality* is needed for explanatory relevance. As Elliott Sober notes, Putnam must think that "the virtue of higher-level explanations does not reside [exclusively] in their greater generality" (1999, 549 n.8). However, Putnam does not explain *why* he thinks this.[9] Sober speculates that Putnam would say that using a physical predicate in an explanation provides "extraneous information," whereas using, say, a psychological predicate does not. Importantly, this extraneous information need not be based on the fact that the psychological predicate is more homotopically abstract (or more general) than the physical predicate. It may instead be grounded in the fact that the psychological predicate is more heterotopically abstract than the physical predicate, *abstracting away from the many physical aspects that are irrelevant to the explanandum at hand*.

---

[9] In fact, Putnam's famous square peg/round hole argument for the relevance advantage seems to contradict the claim in the footnote by appealing to generality. He writes: "the higher level explanation [in terms of rigidity and geometry] is far more general, which is why it is *explanatory*" (1975, 297, italics in original).

5. Multiple Expressibility by Heterotopically Abstract Predicates

Above, I noted that Kitcher appeals to the multiple realizability of PS-processes not only to defend the advantage in explanatory relevance that cytological predicates have over molecular predicates but *also* to argue that this advantage is *non-anthropocentric*. Given that I have argued that the relevance advantage is not dependent on generality and multiple realizability, one might worry that the relevance advantage is now too subjective.

There are a couple related objections here. First, there is the worry that explanatory relevance, in general, is overly anthropocentric. After all, whether or not a predicate is relevant to the explanation of a given phenomenon depends upon what *we* are interested in explaining about that phenomenon, and this will vary over historical time and from culture to culture. I think that this worry can be addressed fairly easily. Of course, our interests are essential for singling out particular phenomena that we find in need of explanation (and even for constructing predicates to express those phenomena). But once those phenomena are singled out, the fact that certain factors are relevant to explaining them, and others are not, is *not* anthropocentric or interest-dependent. For example, given that we are interested in explaining why a man winced and moaned, the intensity and duration of his pain will be very relevant, while the mass and opacity of his nerve fibers will be less so. We can allow that the world has a "built-in" structure of natural phenomena and that which natural phenomena we have access to is partly a function of our cognitive capacities and interests. This is of course compatible with there being natural phenomena that do not align with our interests or that are inaccessible to us.

However, this line of response leads to a subtler objection. Given that we are admitting that there are natural phenomena that are inaccessible to us, the reductionist can claim that there are some *inaccessible* special science phenomena relative to which (possibly disjunctive) physical predicates *do not* contain any irrelevant, extraneous information. If we were smarter, we might have different interests to which "messy" physical predicates, and their associated aspects, were entirely relevant.

Further, the objection continues, these inaccessible explanations might do all the work done by current special science explanations.[10] If this response were successful, it would show that special sciences exist only because we are focusing on an arbitrarily limited set of legitimate explanatory interests and explananda. In principle, physical aspects can do all of the explanatory work that needs to be done.

I think that this objection can also be successfully rebutted. The key point is that any physical property that is plausibly type-identical to some, say, psychological property will not only be characterized *physical* aspects that are irrelevant to some behavioral explanandum, it will also likely be characterized by *other special science* aspects that are irrelevant to that behavioral explanandum. In other words, since a physical property, R, for which type-identity is plausible must be a *total realizer* of a given special science predicate, *M*, it is likely that it will be multiply expressible by several other special science predicates, each of which is heterotopically abstract relative to the physical predicate that expresses R and each of which belongs to an aspect space distinct from the one to which *M* belongs. When this occurs, let us say that R is *heterotopically multiply expressible*.

Examples of heterotopic multiple expressibility appear to be widespread. For instance, consider a sample of gold, which satisfies the predicates 'is rigid,' 'has a high electrical and thermal conductivity', 'is ductile', and 'has a distinctive luster'. These predicates all express a complicated micro-based physical property of the gold sample, the core of which is the cloud of free electrons that permeates the metal (see Menzies 1988 for this example). Or, consider the predicate 'is in pain.' As Colin Allen notes, "pain is processed by multiple [brain] regions, in a highly distributed system; there

---

[10] We need not turn to inaccessible phenomena to make a related point. Sometimes we seek explanatory depth rather than breadth; for example, we may be interested in the microphysical details that explain why this *particular* square peg (made of, say, gold) doesn't fit through this *particular* round hole (see Sober 1999, 548ff.). And quite specific *physical* aspects might be relevant to this question.

is no single brain region whose destruction completely abolishes the experience of pain, sensory processing occurs in areas classically associated with affect, and affect itself is at least a three-layer process with the prefrontal cortex being mainly important for the latest stage of processing" (2004, 620). Pain processing appears to occur in parallel in a plastic, bilateral system activating neurons in the cerebellum, anterior cingulate cortex, insula, thalamus, ventral premotor cortex, and prefrontal cortex, among other areas (Coghill et al. 1999). These findings strongly suggest that if we have a plausible candidate for a micro-based property that is a *total* realizer of 'is in pain', it is likely that this property will also be expressed by distinct psychological predicates concerning affect, motor control, attention, and memory consolidation, some of which may not be conscious and all of which are not easily picked out in everyday language (rough descriptions might include: 'is highly vigilant', 'has narrowly focused attention', and 'is in a certain motivational state'). In both of these cases the realized special science predicates cross-cut one another; they are different kinds of predicates, characterized by different aspects.

Heterotopic multiple expressibility allows us to strengthen the argument that (R2) is false for *non-anthropocentric* reasons, without relying on generality/multiple realizability. Reductive explanation of a psychological generalization like "acute pain usually causes wincing" (to take a crude example) need not fail solely because 'is in pain', is *realized in* a motley of neural or electronic ways (as suggested by Kitcher). Rather, reductive explanation may also fail because a motley assortment of special science predicates each *express* the single complicated physical property, R, that realizes 'is in pain'. Using the physical predicate 'is R' in a derivation of the psychological generalization at best tells us that *some* of the aspects that characterize it (associated with the various the special science predicates that express the property R) are responsible for the truth of this generalization. Thus, the physical explanation amounts to using the *disjunction* of all the special science predicates that express the realizer R: 'is in acute pain' along with, e.g., 'is highly vigilant' and 'has narrowly focused attention'. (For a related point, see Gasper 1992, 668-9.) So, this

physical derivation of the psychological generalization will be objectively deficient as an explanation because it includes all of the irrelevant aspects that characterize these *other* special science predicates. By contrast, citing the predicate 'is in pain' singles out only those aspects that are causally relevant to the wincing.

Similarly, (R3) is false not merely for anthropocentric reasons. Psychological explanations may be deeper and theoretically more fecund than physical explanations because they isolate causal aspects of R that are relevant to some specific behavior (like wincing) from those that are irrelevant. The explanation in terms of 'is R' may contain *more* information since it is expressed by many special science predicates and plays a role in a corresponding number of sets of causal interactions, but not all of this information is relevant to the mental or behavioral explanandum at issue. Physical predicates and aspects are unable to distinguish between the many different sets of special science causal interactions in which the property R participates.

To take another example, even if we are interested only in why, say, some particular square gold peg does not fit through a round hole, the physical predicate that expresses the total physical realizer of 'is rigid' will include irrelevant detail. This physical predicate will also be expressed by special science predicates like 'is ductile' and 'has a distinctive luster,' which are characterized by aspects that are irrelevant to the peg's not fitting through the hole.

A final objection: someone might simply deny that the relevant total physical realizers are expressible by different kinds of special science predicates. That is, one might allege that by singling out certain components of the total micro-based realizer, we can obtain a "stripped down" property that will not be heterotopically multiply expressible—a property that is uniquely expressed by a given special science predicate (and by all of its rough synonyms and more or less *homotopically* abstract predicates). The crucial case to consider is whether a "minimally expressed" physical realizer of say, 'is in pain,' – one for which no "smaller" realizer is also expressed by 'is in pain' – is heterotopically multiply expressible.

Whether many "minimally expressed" realizers are heterotopically multiply expressible is ultimately an empirical question. However, I think that some general considerations about total realizers suggest that many will be. On the one hand, if one strips down the physical realizer solely with respect to its individual *physical* aspects or components in an effort to achieve unique expressibility, then we will no longer have a property expressed by 'is in pain.' For example, singling out aspects like *conductivity* or a property like *having C-fibers firing* will result in a realizer that is not sufficient for the instantiation of *pain*. For, pain requires complex interactions between different *kinds* of physical aspects in a particular context in order to be instantiated. (*Having C-fibers firing* is not sufficient for pain; it is at best a *core* realizer.) On the other hand, if one appeals to "high-level," functional aspects in an attempt to pare down the realizer, then one has effectively conceded that (R2) and (R3) are false. One has accepted that physical aspects are not sufficient for formulating all of the generalizations that there are.

6. Concluding Remarks

Even in the face of the local reduction and disjunctive strategies, some philosophers may want to defend strong NRP instead of accepting (R1) (perhaps by arguing for a hyperintensional criterion of property identity). However, if the arguments in this paper are on the right track, they need not do so in order to offer non-anthropocentric reasons for weak NRP. By grounding the greater explanatory relevance of special science predicates on their heterotopic abstraction relative to physical predicates, defenders of weak NRP can still endorse a claim that all non-reductionists should accept: there are special sciences because of the way the world is put together, and not merely because of the nature of our epistemic relation to the world (cf. Fodor 1974, 113)

References

Allen, Colin. 2004. "Animal Pain." *Nous* 38: 617-643.

Antony, Louise. 1999. "Multiple Realizability, Projectibility, and the Reality of Mental Properties."
    Philosophical Topics 26: 1-24.

Clapp, Lenny. 2001. "Disjunctive Properties: Multiple Realizations." *Journal of Philosophy* 98: 111-
    136.

Coghill, Robert C., et al. 1999. "Pain Intensity Processing Within the Human Brain: A Bilateral,
    Distributed Mechanism." *Journal of Neurophysiology* 82: 1934-1943.

Fodor, Jerry. 1974. "Special Sciences" *Synthese* 28: 97-115.

Gasper, Philip. 1992. "Reduction and Instrumentalism in Genetics." *Philosophy of Science* 59: 655-
    670.

Haug. Matthew C. 2010. "Realization, Determination, and Mechanisms." *Philosophical Studies*. 150:
    313-330.

Kim, Jaegwon. 1998. *Mind in a Physical World*. Cambridge, MA: MIT Press.

———. 1999. "Making Sense of Emergence." *Philosophical Studies* 95: 3-36.

Kitcher, Philip. 1984. "1953 and All That. A Tale of Two Sciences." *Philosophical Review* 93: 335-
    373.

Lewis, David. 1994. "Reduction of Mind." in *A Companion to Philosophy of Mind*. Ed. Samuel
    Guttenplan. Oxford: Oxford UP. Reprinted in David Lewis, *Papers in Metaphysics and
    Epistemology*. Cambridge: Cambridge UP. 8-55.

Menzies, Peter. 1988. "Against Causal Reductionism." *Mind* 97: 551-574.

Putnam, Hilary. (1975) "Philosophy and Our Mental Life." in *Mind, Language and Reality*.
    Cambridge: Cambridge UP. 291-303.

Shapiro, Lawrence A. 2004. *The Mind Incarnate*. Cambridge, MA: MIT Press.

Shoemaker, Sydney. 2007. *Physical Realization*. Oxford: Oxford UP.

Sober, Elliott. 1999. "The Multiple Realizability Argument Against Reductionism." *Philosophy of
    Science* 66: 542-564.