

# Abstraction, Multiple Realizability, and the Explanatory Value of Omitting Irrelevant Details

Matthew C. Haug

The College of William & Mary

[mchaug@wm.edu](mailto:mchaug@wm.edu)

## Abstract

Anti-reductionists hold that special science explanations of some phenomena are objectively better than physical explanations of those phenomena. Prominent defenses of this claim appeal to the multiple realizability of special science properties. I argue that special science explanations can be shown to be better, in one respect, than physical explanations in a way that does not depend on multiple realizability. Namely, I discuss a way in which a special science explanation may be more abstract than a competing physical explanation, even if it is not multiply realizable, and I argue that this kind of abstraction can be used to support the idea that that special science explanation omits explanatorily irrelevant detail.

## 1. Introduction

Almost twenty years ago, Ned Block wrote of a long-standing “anti-reductionist consensus” in the philosophy of mind (and the philosophy of the special sciences more generally), which holds “that reductionism is a mistake and that there are autonomous special sciences” (1997, 107). Although this consensus seems to have been weakened in recent years, anti-reductionism is still arguably the dominant view. Disagreement remains about exactly what it is for a special science to be autonomous, but I take it that autonomy involves at least the following claim:

*Explanatory autonomy*: Some special science explanation of a given fact or event is objectively better than any fundamental physical explanation of that fact or event.<sup>1</sup>

The classic papers that defended *explanatory autonomy*, and instituted the anti-reductionist consensus of which it is a part (e.g., Putnam (1967, 1975), Fodor (1974), and Kitcher (1984)), appealed to the alleged *multiple realizability* of special science properties. That is, these papers argued that any given special science

---

<sup>1</sup> *Explanatory autonomy* pits the explanations offered by fundamental physical theory against all other scientific explanations. In practice, however, debates about reduction have often been waged between explanations offered by a pair of “nearby” sciences, such as classical genetics and molecular genetics (e.g., Kitcher 1984) or cognitive psychology and neuroscience (e.g., Bechtel and Mundale 1999). Two common assumptions support this practice: first, that sciences are organized into hierarchical levels and, second, that sciences at the “lower levels” of this hierarchy (such as molecular genetics and neuroscience) are reducible to fundamental physics. Anti-reductionists and reductionists, then, disagree about whether sciences at higher levels in the hierarchy are reducible to relatively lower level sciences, with anti-reductionists claiming that reducibility fails to hold at some point in the hierarchy. Although there are reasons to doubt whether these assumptions are true, I shall adopt them in this paper for the sake of argument, along with the claim that neuroscience is reducible to fundamental physics. If these assumptions are false, then either anti-reductionism is in an even better position or the debate about reduction itself needs to be re-framed.

property could be realized by different (physically heterogeneous) physical properties and that this multiple realizability was crucial for defending *explanatory autonomy*.<sup>2</sup> In this paper, I argue that there is at least one respect in which some special science explanations are better than physical explanations that need not be “based on an argument from multiple realizability.” Namely, some special science explanations isolate explanatorily relevant features even if they do not involve properties that are multiply realizable. As I discuss below, these explanations are more *abstract* than competing physical explanations in a way that does not entail that they are more general. If I am right, it is a mistake to think that *explanatory autonomy* hinges on multiple realizability. Avoiding this mistake sidesteps the morass of debates about the nature and extent of multiple realizability and opens up new conceptual space in an old debate.

## ***2. Relevance, Generality, and Multiple Realizability***

One plausible way in which an explanation is objectively good is if it provides the right amount of explanatory detail (see, e.g., Garfinkel 1981; Batterman 2002; Strevens 2008). For instance, suppose that an extremely loud utterance of the word “shatter” causes a crystal wine glass to break (cf. Dretske 1988, 79). An explanation that cites the pitch and intensity of the sound waves is better than one that also cites the meaning of the utterance. This is because the meaning of the utterance was irrelevant to the glass’s breaking. The glass would have broken regardless of the utterance’s meaning (or lack thereof). Examples like this one support the following explanatory virtue:

*Relevance.* Other things being equal, one explanation is better than another if it includes fewer features that are irrelevant to (the production of) the fact or event to be explained.

One explanation will be better, with respect to *relevance*, than another to the extent that it is more *abstract*, in the sense that it eliminates or omits details.

There are other good-making features of explanations aside from *relevance*. For instance, good explanations will unify disparate phenomena. Other things being equal, one explanation is better than

---

<sup>2</sup> As Block notes, it is part of the consensus that the autonomy of the special sciences is “based on an argument from multiple realizability” (1997, 107).

another if it is more *general*, if it is, in some sense, more widely applicable. One kind of generality that has often been used to defend *explanatory autonomy* traces back to Putnam's (1967, 1975) work: one explanation of a given fact or event is more general, in this sense, than another if applies to wider range of (logically) possible situations. As with *relevance*, explanations will be more general to the extent that the properties that figure in them are more abstract.<sup>3</sup> By eliminating or omitting physical details, some abstract special science properties will apply to more possible situations.

Since both the greater relevance and greater generality of special science explanations depend on their being more abstract than physical explanations, it can seem as if these two explanatory virtues stand or fall together. If there were a single kind of abstraction that was required to support both of these virtues, then this would go a long way toward explaining why influential defenses of the greater *relevance* of special science explanations (and not just their greater *generality*) have relied on multiple realizability. Here I briefly sketch two such defenses.

First, consider Putnam's famous discussion of why a 15/16" rigid cubical peg fits through a 1" square hole, but not a 1" circular hole, in a rigid board. Putnam claims that an explanation of these facts in terms of the shapes and relative sizes of the peg and holes and the rigidity of the peg and board "brings out" the "*relevant structural features of the situation*," which any microphysical explanation "conceals" (1975, 296, 297, italics in original). And he suggests that this is because the "higher level" structural features are multiply realizable: the same explanation in terms of these features will be correct "whether the peg consists of molecules, or continuous rigid substance, or whatever" (ibid., 296). The idea is that changing the underlying microstructure makes no difference to whether the square peg goes through the round hole; all that is relevant is the abstract structural features that are *common* to those different physical realizers.

Similarly, Kitcher (1984) 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

---

<sup>3</sup> In fact, in a paper that isolates and defends this kind of generality as a dimension of explanatory depth, Brad Weslake (2010) simply calls it *abstraction*.

rearrangements the derivation would only blur the outline of a simple cytological story, *adding a welter of irrelevant detail*” (ibid., 347, italics added). A little later, he claims that “adding [molecular] details would only disguise the relevant factor,” namely, that meiosis is a “pair-separation process” (ibid., 348). And, he claims that this fact about relevance hinges on multiple realizability—on the alleged fact that pair-separation processes are “heterogeneous from the molecular point of view” (1984, 349) because they are “realized in a motley of molecular ways” (ibid., 350).<sup>4</sup>

Basing the greater *relevance* of special science explanations on multiple realizability may be ill advised, however. For, almost contemporaneously with the formation of the anti-reductionist consensus, some philosophers raised serious doubts about the multiple realizability argument,<sup>5</sup> and, in the two decades since Block’s (1997) article appeared (in which he responded to one of these lines of doubt) a number of authors have raised additional critiques concerning the extent of multiple realizability and whether it can be used to establish the explanatory superiority of the special sciences.<sup>6</sup> For all I say in this paper, the multiple realizability argument may be successfully defended against such critiques. However, if I am right, such a defense is unnecessary: the greater *relevance* of some special science explanations can be supported by a kind of abstraction that need not involve multiple realizability.

### *3. A Distinction Among Kinds of Abstraction*

One entity is more abstract than another if it lacks detail that the other possesses. Different instances of abstraction will involve omission or elimination of different kinds of details, some of which I discuss below. There is an important distinction to be drawn among these varieties of abstraction: some kinds of abstraction require that more abstract entities are more *general* than less abstract ones (call this *generality-entailing (GE) abstraction*), while other kinds of abstraction do not impose this requirement

---

<sup>4</sup> See also Garfinkel (1981), who claims that a “microexplanation” of why a rabbit was eaten (in terms of the “equations of interaction between individual foxes and individual rabbits (depending on such things as their physiology and reaction times)” (ibid., 55)) is inferior to a “macroexplanation” that simply cites the high population-density of foxes in the area (and implicitly, the Lotka-Volterra equation). Garfinkel claims that this is because the microexplanation “contains much that is *irrelevant* to why the rabbit got eaten and ... [these irrelevant data] *bury* the explanation unrecognizably” (ibid., 56, italics added). And *this* is true, according to Garfinkel, because the macroexplanation is *stable* through variation in the way that the high population-density of foxes is realized (ibid., 57).

<sup>5</sup> See, e.g., Lewis (1969), Kim (1972, 188-92).

<sup>6</sup> See, e.g., Bechtel and Mundale (1999); Shapiro (2004)

(call this *non-generality-entailing (NGE) abstraction*). I first discuss two varieties of abstraction that fall on the former side of this divide; then, I outline two kinds of abstraction that fall on the latter side.

First, consider the way in which determinables are more abstract than their determinates. *Being red* is more abstract than *being crimson* or *being scarlet*. What kind of detail does a determinable lack that its determinates possess? Assuming that colors are individuated by three “determination dimensions”: hue, saturation, and brightness, *being crimson* will be associated with particular values (or a small range of values) of hue, saturation, and brightness. It will occupy a relatively small portion of a three-dimensional color space. *Being red*, by contrast, will not be specified at this fine-grained level of detail. Rather, it will be characterized by a broader range of hue, saturation, and brightness values. It will occupy a larger portion of three-dimensional color space, a portion that includes the space occupied by *being crimson* as a proper subspace. Following Haug (2011), I’ll call this kind of abstraction, *homotopic* abstraction, since it applies to properties that are characterized by the *same property space*. In general, we can say that a property P is more homotopically abstract than a property Q if and only if Q occupies a proper subspace of the portion of the property space occupied by P.<sup>7</sup>

Second, consider the relation between a multiply realizable property and each of its individual realizers. There is disagreement about how to characterize multiple realizability, but on any plausible account, multiply realizable properties are more abstract than their individual realizers. If property P (of type X) is multiple realizable, then P omits (or “abstracts away”) from features that are unique to its individual realizers and isolates features that these realizers (or the objects that possess them) have in common. This will be true whether these features are causal powers (Wilson 1999; Shoemaker 2001; Gillett 2002), or functional roles (Shapiro 2004; Polger 2007), or exact similarity of X-type features (Funkhouser 2007).

---

<sup>7</sup> Note that this differs somewhat from the definition given by Haug (2011). Eric Funkhouser (2006, 2014) provides a helpful discussion of “determination dimensions” and how they (together with “non-determinable necessities”) specify the nature of kinds of properties. Note that homotopic abstraction is the inverse relation of what Funkhouser calls “specification.”

Homotopic abstraction and multiple realizability both clearly require more abstract properties to be more general than less abstract ones: they are instances of GE-abstraction. Clearly, the property *being red* must apply not just to scarlet objects but also to crimson ones. Similarly, a multiple realizable property, P, must apply not only to objects that possess one of P's realizers, R, but also to objects that possess another of P's realizers (and do not have R).

Now consider a third example of abstraction. Given a physical system, one model or set of equations describing that system will be more abstract than another if it eliminates or omits at least one feature of that system (e.g., a quantity, degree of freedom, or boundary or initial condition) that the latter model or set of equations includes. For example, suppose we have a rotating sphere moving with a constant linear velocity along the x-axis. A set of equations ( $\mathcal{p}$ ) that includes only an equation of motion for the sphere's linear motion is more abstract than one ( $\mathcal{p}L$ ) that includes equations of motion for *both* its linear motion *and its rotation*. Similarly, a thermodynamic model of a gas that takes the number of particles and volume to be infinite (while maintaining a constant ratio of number of particles to volume) is more abstract than a model that includes the boundary conditions imposed by the gas's container. (For these examples, see Knox (2016, 44-45, 50).)

Finally, consider the relation between *having some hue or other* (i.e., *being hued*) and *having some color or other* (i.e., *being colored*). *Being hued* is more abstract than *being colored*; it omits the other two dimensions of color: saturation and brightness. Following Haug (2011), we can call this kind of abstraction, *heterotopic abstraction*, since it applies to properties from *different* property spaces. If a property Q metaphysically necessitates a property P, then P is more heterotopically abstract than Q if and only if the characteristic property space of P has fewer dimensions than the characteristic property space of Q.<sup>8</sup>

These two kinds of abstraction do *not* require that more abstract entities are more general than less abstract ones. A more abstract model or set of equations need not apply to more systems than a less abstract model or set of equations. For example, suppose that we take the set of equations  $\mathcal{p}L$  above and

---

<sup>8</sup> Note that this also differs slightly from the account given by Haug (2011).

stipulate that the linear momentum is zero; call this set of equations (including the trivial equation of motion for rotation)  $pL\text{-zero}$ . Then, the set of equations  $p$  (that includes only an equation for linear motion) will apply (i.e. accurately describe) the kinematics of *exactly the same* set of possible physical systems as  $pL\text{-zero}$ , namely, those spheres moving along the x-axis that have zero angular momentum. However,  $p$  is still more abstract than  $pL\text{-zero}$  since it *omits* the equation of motion for rotation entirely.<sup>9</sup>

Similarly, unlike the pair of *being red* and *being scarlet*, *being hued* is not more general than *being colored*. Necessarily, anything that has a hue is also colored (and, necessarily, anything that is colored is also hued). These properties are necessarily co-extensive, even though one is more abstract than another. Thus, in these cases, greater abstraction does not require greater generality: they are instances of NGE-abstraction.

#### **4. NGE-abstraction and Explanatory Relevance**

With the distinction between GE-abstraction and NGE-abstraction in hand, we can see that the former (and thus multiple realizability) is not required for defending the greater *relevance* of special science explanations; NGE-abstraction will work at least as well.

First, suppose that we have non-rotating sphere moving along the x-axis with a constant velocity. The systems of equations  $pL\text{-zero}$  and  $p$ , from above, apply to exactly the same systems;  $p$  is not more general than  $pL\text{-zero}$ ; nevertheless,  $p$  clearly provides a better explanation of the sphere's position; it isolates only the features that are relevant to this explanandum, omitting the irrelevant (and, incidentally, zero-valued) angular momentum. (For a more substantive example of NGE-abstraction with respect to certain features of a system, see Section 5.)

Now, return to Putnam's example from Section 2, and note that the "high level" structural properties in Putnam's case are more heterotopically abstract than the microphysical properties that underlie them. *Being a circle* and *being a square* each have a single determination dimension (diameter length and side length, respectively). Further, *rigidity* (also known as *stiffness*) also has a single determination

---

<sup>9</sup> This example is inspired by Knox's (2016, 52) discussion of idealization as a "precursor to abstraction."

dimension measured in units of [force/distance]. By contrast, the microphysical properties that are possible realizers of these structural features are characterized by more determination dimensions, including at least: (1) the nature of the components that make up the peg and board, (2) the nature of the bonds that hold between these components (e.g., ionic, covalent, metallic, or none (in the case of a (logically possible) continuous rigid substance)); (3) the angle(s) between these bonds; and (4) the entropy of the peg and board (which is especially important for understanding the stiffness of polymers such as rubber). (See Roylance (2000) for a discussion of the microscopic basis of *stiffness*.)

A correct explanation of the fact that the square peg does not pass through the round hole depends only on the relations between the determination dimensions of *being a circle*, *being a square*, and *rigidity*. The determination dimensions of the underlying realizers are irrelevant to this explanation. Importantly, it is not the *multiplicity* of the *values* of these other determination dimensions that is crucial here but simply the fact that these other dimensions characterize the microphysical realizers *at all*. That is, it is not facts about homotopic abstraction or multiple realizability that are crucial for *relevance* but rather facts about heterotopic abstraction, a kind of *non-generality-entailing* abstraction.<sup>10</sup>

One might worry that multiple realizability (or homotopic abstraction) is required to justify the claim that these other dimensions are in fact explanatorily *irrelevant* to the explanandum—that without showing that the explanatory relationship is “stable” or “robust” through *variation* in the *values* of these underlying dimensions we would have no evidence that these other dimensions are *themselves* explanatorily irrelevant. However, this is not the case. Even without showing that the generalization that 15/16” square pegs do not fit through 1” circular holes is independent of variation in the underlying physical realizers, it remains the case that this generalization is realization independent in the sense that one can discover and confirm it without knowing anything about the physical realizers that underlie it (much less confirming a generalization between them). That is, the very fact that structural predicates like

---

<sup>10</sup> Haug (2011) claims that heterotopic abstraction can support the idea that some special science explanations omit explanatorily irrelevant details that are included in physical explanations. However, the discussion of this point is very brief, and Haug does not discuss how other forms of NGE-abstraction can also support this idea. (Cf. the discussion of the sphere example above and of thermodynamics in Section 5.)



“is a 15/16” square peg” are projectible is strong presumptive evidence that such predicates pick out objective features of the world that figure in real regularities. (On this point, see Antony (1999, 14ff.)) The underlying microstructure (whether it maps many-to-one or one-to-one to macrostructure) is irrelevant to confirming that such regularities obtain.

We can get clearer about how heterotopic abstraction can be used to defend the greater *relevance* of special science explanations without appealing to multiple realizability by seeing how such a defense effectively blocks reductionist appeals to “disjunctive properties” or “local reductions” (see, e.g., Kim 1992; Lewis 1994). These appeals, in effect, claim that (perhaps disjunctive) physical properties can in fact be matched up one-to-one with (at least structure- or species-restricted) special science properties and that explanations in terms of these physical properties will be just as *relevant* and *general* as (any genuine) special science explanations. However, while these reductionist gambits may be successful with respect to *generality*, facts about heterotopic abstraction show that they are ineffective at undermining the greater *relevance* of special science explanations.

Consider a disjunctive property, D, that has every metaphysically possible realizer of *being a 15/16” square peg* as a disjunct. D is just as general as *being a 15/16” square peg*; they are necessarily co-extensive. Thus, Putnam’s multiple realizability argument for the greater *relevance* of the macro-structural explanation is undermined.<sup>11</sup> However, D’s instantiation consists merely in one of its disjuncts being instantiated, so its property space is arguably the *sum* (i.e. the span of the union) of the property spaces of each of its disjuncts. That is, D’s set of determination dimensions is the union of the sets of determination dimensions of its disjuncts. Thus, *being a 15/16” square peg* is still more heterotopically abstract than D; it still isolates features that are explanatorily relevant to the fact that the peg fails to pass through the round hole, while omitting features that are irrelevant to this behavior.

---

<sup>11</sup> Implicit in the above discussion is Putnam’s claim that “the higher level explanation is far more general [than the microphysical one], which is why it is *explanatory*” (1975, 297, italics in original). But this claim is not true when the microphysical explanation is in terms of D.

Funkhouser also responds to the “disjunctive property” objection to the multiple realizability argument by pointing out that a realized property R has different determination dimensions than a disjunctive property whose disjuncts are each of R’s possible realizers (2014, 108-9). But Alexander Bird (2015) rightly asks how this point alone constitutes a reply to the objection. The key point that Funkhouser leaves out is that a difference in determination dimensions—in particular, greater heterotopic abstraction—supports greater explanatory *relevance* even if the absence of multiple realizability. Funkhouser emphasizes that realized properties and their realizers, by having different determination dimensions, are at different “levels of abstraction” (2014, 78, 89, 124). However, he does not seem to recognize fully that (as the disjunctive realizer case illustrates) this kind of abstraction does *not* depend on *multiple* realizability but rather on the realization relation (a particular kind of asymmetric necessitation) itself.

### ***5. Implications for How to Frame Debates about Reduction and Autonomy***

In the last few years, several philosophers have defended explanatory autonomy or irreducibility without relying on multiple realizability (e.g., Wilson 2010; Knox 2016). I’ll conclude by briefly applying the notion of NGE-abstraction to one of these defenses (Knox 2016) and suggesting that this can help us better understand how to formulate debates about explanatory autonomy and reducibility, in general.<sup>12</sup>

Knox (2016) argues that thermodynamics offers novel explanations of many phenomena, such as why diesel engines, unlike gasoline engines, do not need spark plugs. The novelty of this thermodynamical explanation consists in the fact it involves an abstraction (namely, omitting all details related to heat transfer) that “cuts across” the quantities that are recognized as natural by statistical mechanics (ibid., 46, 56). Further, Knox claims that this fact about explanatory novelty is compatible with the *reducibility* of thermodynamics to statistical mechanics. According to Knox, the key to

---

<sup>12</sup> Wilson’s (2010) argument that “eliminations of degrees of freedom” are sufficient for ontological irreducibility also relies on NGE-abstraction and not multiple realizability, but a discussion of this fact (and its implications) will have to await another occasion.

understanding this compatibility is that the bridge laws linking thermodynamical and statistical mechanical quantities will involve complex, “mathematically irreversible” operations (such as taking the limit as the number of gas particles goes to infinity) and thus will themselves involve abstraction (ibid., 54, 57). As a result, further abstractions with respect to thermodynamical quantities will be “opaque” from the perspective of statistical mechanics (ibid., 42).

First, note that these further abstractions will be instances of the first kind of NGE-abstraction discussed above—ignoring heat transfer simplifies the thermodynamical equations and isolates the explanatorily relevant features. Knox claims that this results in “novel explanations that are not merely abstractions of some more detailed [statistical mechanical] picture” (ibid., 41), and she seems to ground this novelty in the “mathematically irreversible” change in variables that occurs when one moves from statistical mechanical to thermodynamical quantities (ibid., 56). However, not every mathematically irreversible operation seems to induce this kind of “unnaturalness” from a lower level perspective: for example, taking a sum or a mean will lead to loss of information, but the result *will* be “merely an abstraction from some more detailed underlying picture.” I think that the second kind of NGE-abstraction—heterotopic abstraction—can help here. It is not mere “mathematical irreversibility” that supports novelty but *heterotopic abstraction*. The mean of a quantity of a system is more homotopically abstract than a description that specifies the particular values of that quantity had by the components of that system: it omits detail within a *single* property space. But moving from a statistical mechanical quantity to a thermodynamical quantity involves moving to a *new* property space, one whose dimensions “cut across” the dimensions of the property space of the statistical mechanical quantity.

Knox notes that her account of explanatory novelty fits poorly into standard taxonomies of emergence/reduction that characterize weak emergence as merely epistemic and strong emergence as metaphysical (2016, 58). Her account is weaker than standard epistemic accounts in that it is compatible with the theoretical reduction of thermodynamics to statistical mechanics, but it is stronger than

epistemic accounts in that it depends on objective features of the world and not merely on our cognitive limitations (ibid., 58, 44).

But how can there be any *objectively* better “high level” explanations if all of the properties involved in those explanations can be mapped one-to-one via *bridge laws* to “low level” properties? What “objective features of the world” could these explanations be tracking other than those of fundamental physics? I think that NGE-abstraction suggests a framework within which to answer these questions. A property can be more NGE-abstract than another with which it is necessarily co-extensive. If NGE-abstraction is itself an “objective feature of the world,” this suggests that we should adopt a *hyperintensional* criterion for property individuation.<sup>13</sup> In short, there are “objective features of the world” that are more fine-grained than, and thus cannot be captured by, the resources used in standard formulations of the metaphysics of reduction and autonomy. Working out this hyperintensional account of property individuation will not be a trivial task, but if the discussion in this paper is on the right track, it is a task that is important not only for defending the autonomy of the special sciences but also for spelling out exactly what such autonomy amounts to.

---

<sup>13</sup> This provides a further motivation for adopting a program of “hyperintensional metaphysics” (Nolan 2014), one that is more closely tied to actual scientific practice than recent work on the hyperintensional notion of “grounding” in analytic metaphysics.

## References

- Antony, Louise. (1999) "Multiple Realizability, Projectibility, and the Reality of Mental Properties." *Philosophical Topics*. 26: 1-24.
- Batterman, Robert W. (2002) *The Devil in the Details*. Oxford: Oxford UP.
- Bechtel, William and Jennifer Mundale. (1999) "Multiple Realizability Revisited: Linking Cognitive and Neural States." *Philosophy of Science*. 66: 175-207.
- Bird, Alexander. (2015) "Review of *The Logical Structure of Kinds*." *Notre Dame Philosophical Reviews*. 23 Oct. <https://ndpr.nd.edu/news/61948-the-logical-structure-of-kinds/>
- Block, Ned. (1997) "Anti-reductionism Slaps Back." *Philosophical Perspectives*. 11: 107-132.
- Dretske, Fred. (1988) *Explaining Behavior: Reasons in a World of Causes*. Cambridge MA: MIT Press.
- Fodor, Jerry. (1974) "Special Sciences." *Synthese*. 28: 97-115.
- Funkhouser, Eric. (2006) "The Determinable-Determinate Relation." *Noûs*. 40: 548-569.
- . (2007) "A Liberal Conception of Multiple Realizability." *Philosophical Studies*. 132: 467-494.
- . (2014) *The Logical Structure of Kinds*. Oxford: Oxford UP.
- Garfinkel, Alan. (1981) *Forms of Explanation*. New Haven, CT: Yale UP.
- Gillett, Carl. (2002) "The Dimensions of Realization: A Critique of the Standard View." *Analysis*. 62: 316-323.
- Haug, Matthew C. (2011) "Abstraction and Explanatory Relevance; or, Why Do the Special Sciences Exist?" *Philosophy of Science*. 78: 1143-1155.
- Kim, Jaegwon. (1972) "Phenomenal Properties, Psychophysical Laws, and the Identity Theory." *The Monist*. 56: 177-192.
- . (1992) "Multiple Realizability and the Metaphysics of Reduction." *Philosophy and Phenomenological Research*. 52: 1-26.
- Kitcher, Philip. (1984) "1953 and All That." *Philosophical Review*. 93: 335-373.

- Knox, Eleanor. (2016) "Abstraction and Its Limits: Finding Space for Novel Explanation." *Noûs*. 50: 41-60.
- Lewis, David. (1969) "Review of *Art, Mind, and Religion*." *Journal of Philosophy*. 66: 23-35.
- . (1994) "Reduction of Mind." in *Companion to the Philosophy of Mind*. Ed. Samuel Guttenplan. Cambridge: Blackwell. 412-431.
- Nolan, Daniel. (2014) "Hyperintensional Metaphysics." *Philosophical Studies*. 171: 149-160.
- Polger, Thomas. (2007) "Realization and the Metaphysics of Mind." *Australasian Journal of Philosophy*. 85: 233-259.
- Putnam, Hilary. (1967) "Psychological Predicates." In *Art, Mind, and Religion*. Eds. W.H. Capitan and D.D. Merrill. Pittsburgh: University of Pittsburgh Press, 37-48. Reprinted as "The Nature of Mental States" in his *Mind, Language, and Reality: Philosophical Papers, Vol. 2*. Cambridge: Cambridge UP. 429-440.
- . (1975) "Philosophy and Our Mental Life." In his *Mind, Language, and Reality: Philosophical Papers, Vol. 2*. Cambridge: Cambridge UP. 291-303.
- Roylance, David. (2000) "Atomistic Basis of Elasticity." *Mechanics of Materials*. (MIT OpenCourseWare) [http://ocw.mit.edu/courses/materials-science-and-engineering/3-11-mechanics-of-materials-fall-1999/modules/elas\\_2.pdf](http://ocw.mit.edu/courses/materials-science-and-engineering/3-11-mechanics-of-materials-fall-1999/modules/elas_2.pdf)
- Shapiro, Lawrence. (2004) *The Mind Incarnate*. Cambridge, MA: MIT Press.
- Shoemaker, Sydney. (2001) "Realization and Mental Causation." In *Physicalism and Its Discontents*. Ed. Carl Gillett and Barry M. Loewer. Cambridge: Cambridge UP. 74-98.
- Strevens, Michael. (2008) *Depth: An Account of Scientific Explanation*. Cambridge, MA: Harvard UP.
- Weslake, Brad. (2010) "Explanatory Depth." *Philosophy of Science*. 77: 273-294.
- Wilson, Jessica. (1999) "How Superduper Does a Physicalist Supervenience Need to Be?" *Philosophical Quarterly*. 49: 33-52.

———. (2010) “Non-Reductive Physicalism and Degrees of Freedom.” *British Journal for the Philosophy of Science*. 61: 279-311.