

Multiple Realizability from a Causal Perspective

Lauren N. Ross

This paper examines the multiple realizability thesis within a causal framework. The beginnings of this framework are found in Sober (1999), which argues that the multiple realizability thesis poses no challenge to reductive explanation. While Sober's causal approach has the potential to reveal new insights, I argue that his set up fails to capture important aspects of the multiple realizability thesis. After correcting for these issues, I argue that this causal framework reveals something quite different. It reveals how multiple realizability relates to a common type of causal complexity in biology that poses problems for reductive explanation.

1 Introduction. The multiple realizability thesis developed by Putnam (1967; 1975) and Fodor (1968; 1974) has received significant attention for its argument against reductive explanation. A key feature of this thesis is its focus on cases where a particular higher-level phenomenon, such as “money” or “pain,” is multiply realized by different lower-level details. These details might consist of different currencies (wampum, dollars, checks, etc.) in the case of money and different physical brain states in the case of pain. In these examples, the one-to-many relationship between predicates in some higher-level science and those in a lower-level science is said to render reductive explanation unattainable. The splintering of predicates at lower-levels prevents them from capturing the single higher-level predicate that they all realize and the causal regularities that this higher-level predicate figures in. These points have been used to suggest that higher-level detail can provide explanations that are “objectively” superior to explanations provided by lower-level detail. This is because some higher-level details capture a kind of generality, unity, and broad scope that is lost at the lower-levels. Such claims have been used to argue for the “autonomy” of the special sciences from lower-level sciences such as physics (Putnam 1975, 131-132).

In an influential paper, Sober (1999) argues against these conclusions of the multiple realizability thesis. He examines this thesis in the context of causal explanation in biology and claims that there are no convincing reasons to view higher-level details as more explanatorily relevant than lower-level ones. The level of detail cited by scientists has to do with the amount of detail that they are interested in, where this is driven by “matter[s] of taste” (Sober 1999, 551). Sometimes they are interested in more detail, sometimes they are interested in less, and each choice is guided by different explanatory virtues. In contrast to what details interests scientists, there are further questions about what details provide “objectively” superior explanations. According to Sober, lower-level details can provide objectively superior explanations compared to higher-level ones. He suggests that for any explanatory target, lower-level details can always be included without detracting from the explanation. The worst offense committed by this extra lower-level detail is that it “explains too much,” while the same cannot be said for higher-level detail (Sober 1999, 547). According to Sober, the multiple realizability thesis has overlooked this consideration and failed to appreciate

*I would like to thank Elliott Sober for helpful feedback on this paper.

†To contact the author, please write to: Lauren N. Ross, Department of Logic and Philosophy of Science, University of California, Irvine, 3151 Social Science Plaza A, University of California, Irvine 92697-5100; email: rossl@uci.edu.

that just because someone “may not want to hear the gory details...does not mean that the details are not explanatory” (Sober 1999, 549). His main points are that the multiple realizability thesis poses no threat to reductive explanation, that a reductive claim actually follows from this thesis, and that various successes in biology are a result of the fact that this multiple realizability thesis has not “won the hearts” of scientists in this domain (Sober 1999, 556). Sober’s argument has received significant praise and attention in the field, with some claiming that his analysis has “definitively refuted” the multiple realizability argument against reduction (Butterfield 2011a, 942) (Butterfield 2011b, 104).

One unique feature of Sober’s (1999) paper is that it analyses causal explanation to support various claims about the multiple realizability thesis and reductive explanation. This approach has a number of potential advantages. It has the potential to capture new challenges that multiple realizability might pose for explanation, how these relate to actual causal scenarios in biology (if at all), and how reductive explanation fares in light of them. However, while this approach has the potential to reveal new insights, Sober’s set up fails to capture important aspects of the multiple realizability thesis.¹ In particular, his analysis examines a token-level explanatory target, while the main strength of the multiple realizability thesis pertains to type-level explanatory targets. After revising Sober’s causal framework to properly capture this feature, I argue that this causal framework reveals something quite different. First, it captures how multiple realizability relates to a common type of causal complexity in biology, which I call causal heterogeneity. Second, it clarifies how multiple realizability and causal heterogeneity do pose problems for reductive explanation. These problems have to do with situations where lower-level detail is less explanatory, in an objective sense, than higher-level detail. In these cases, explanatory power has to do with the scope of a factor’s control over instances of a type-level explanatory target. This notion of scope is related to how “proportional” a cause is to its effect, or how well it “fits” with its effect (Yablo 1992; Woodward 2010). Third, I suggest that while these cases do reveal challenges for reductive explanation they do not support the strict “autonomy” of higher-level phenomena from lower-levels. This analysis supports a more measured view about the consequences of multiple realizability for reductive explanation. Some higher-level biological explanations enjoy a kind of “autonomy” from lower-level details and others have a reductive character. The particular level of detail that is cited in these explanations is determined by the level at which relevant causal factors are located and these factors may be located anywhere from (i) lower (ii) mid-range, to (iii) higher-levels.

2 Multiple realizability. Discussions of reductive explanation and multiple realizability often start with a layered view of the world in which it is comprised of hierarchical levels. The lowest level is occupied by physics, the next highest level by chemistry, and further increasing levels include biology, psychology, and various social sciences (Wimsatt 1976; Sober 1999). Given this picture,

¹My analysis examines multiple realizability and reductive explanation in the context of actual biological science. This differs from the approach found in Sober’s work. Sober is primarily interested in explanatory reduction “in principle” as opposed to “in practice” (Sober 1999, 543). He is concerned with the possibility of explanatory reduction for “future science” or some “complete physics” even if this is not yet found in current science (Sober 1999, 543). Despite these differences, Sober still considers biological case studies, scientific reasoning in biology, and the influence of these debates on scientific practice. However, I will refrain from extending my analysis to claims about “in principle,” “future,” or “complete” science. Basing an analysis off of “future” science seems contentious and unhelpful—we do not know what the future will hold and these guesses about the future do not represent actual science. I aim to show that multiple realizability and causal heterogeneity pose problems for reductive explanation in actual biological science.

the basic idea of reduction is that “all the special sciences reduce to physics” (Fodor 1974, 98). This is often understood as the claim that the laws or causal generalizations figuring in explanations of the higher-level sciences can be reduced to laws or causal generalizations that operate at the lower-levels of physics. While reductionism is ultimately concerned with reducing these higher-level regularities down to the level of physics, many analyses focus on the possibility of reduction from a given level down to *any* lower-level. This larger reductive project is thought to depend on and involve these more local cases of reduction.

2.1 Putnam and Fodor. In order to explore this further, consider the higher-level regularity “smoking causes lung cancer.” This can be captured by “ $P \rightarrow Q$ ” in Figure 1 where P = smoking, Q = lung cancer, and the arrow represents their causal connection. This captures a higher-level causal explanation in the sense that causes explain their effects. Standard accounts claim that reducing this higher-level explanation to a lower-level one involves at least two main things. First, (i) predicates in the higher-level explanation should be rendered into predicates familiar to the vocabulary of the lower-level science. In this step, “bridge principles” connect terms in the vocabularies of each science. These connections are represented by the dotted lines in Figure 1. These connections are “symmetric” identity relations in the sense that they depict higher-level properties and those lower-level details that instantiate or realize them. For example, bridge principles connect the higher-level predicate “smoking” (P) to the particular lower-level carcinogens (A_1 , A_2 , or A_n) that realize or instantiate this property. In the second step, (ii) the laws or causal regularities that involve these higher-level predicates should be translated into laws or regularities in the lower-level science. These laws or regularities are asymmetric *causal* relations, as opposed to the symmetric *identity* relations discussed in the first step (Fodor 1974, 99). Such causal relations are represented by the horizontal arrows in Figure 1, which capture the direction of causal influence.




Figure 1: Redrawn from Fodor (1974, 109).

According to Fodor and Putnam, there is at least one main problem for this reductive project. The problem is that phenomena in the higher-level science are multiply realized by phenomena in the lower-level science in a way that prevents the lower-level science from capturing desirable features of the higher-level explanation. In particular, the lower-level science does not capture the single higher-level regularity that the distinct lower-level properties all figure in or what it is that these lower-level properties all have in common. In order to see this, consider the fact that the higher-level explanation unifies a class of events that share important features. The higher-level explanation captures how all (or most) cases of smoking result in a significantly increased risk of

acquiring lung cancer no matter which of a number of different carcinogens are involved. The similarity of these cases cannot be captured at the lower-level science because they are “wildly disjunctive” at this level (Fodor 1974, 103). In other words, the higher-level regularity captures an explanatory pattern that is lost at the lower-level because it is “about events whose [lower-level] physical descriptions have nothing in common” (Fodor 1974, 103). In this sense, the lower-level science “conceals” (Putnam 1975, 132) common features across these cases and, in doing so, it fails to capture their “epistemologically important properties” (Fodor 1974, 103). These properties are captured in the higher-level science, which is said to provide a superior explanation because of its broad scope, application to a wide range of systems, and the fact that it can generalize to other cases (Putnam 1975). The inability to capture these features at the lower-level marks a failure of the reductive project. This is not to say that the lower-level details are entirely unexplanatory, but if they are explanatory they provide a “terrible” explanation (Putnam 1975, 131). Fodor and Putnam are careful to clarify that this multiple realizability thesis is not problematic for token-level physicalism (which is often confused with reduction) or the claim that *single instances* of higher-level phenomena are instantiated by particular lower-level details (Fodor 1974, 105). The proper target of the multiple realizability thesis is reductivism, which entails a kind of *type-level* physicalism in which type-level predicates in a special science correspond to type-level predicates in a lower-level science. In other words, while one might be able to “reduce” a single instance (token-level) of some outcome to a lower-level science, the claim is that this cannot be done for multiple instances of a particular higher-level outcome (type-level).² Failure of the later (type-level physicalism) is the focus of the multiple realizability thesis.

2.2 Sober’s analysis. Sober frames his analysis of the multiple realizability thesis with Figure 2. This figure illustrates a higher-level causal relationship where a single instance of x causes a single instance of y . While the higher-level property x can be instantiated by a variety of lower-level details, the particular instance of x in question is instantiated by the particular lower-level properties z . The symmetric relationship of identity (or realization) between x and z is represented by arrow “d,” while the asymmetric relationships of causation from x to y , and from z to y are represented by arrows “e”.³ Sober focuses his analysis on the latter causal relationships because “these provide the clearest cases of scientific explanations” (Sober 1999, 546). He claims that the multiple realizability thesis bears on whether the higher or lower-level causes are equally explanatorily relevant to y and that this thesis views the higher-level cause as providing a better explanation. According to Sober, the multiple realizability thesis maintains that “(z) does not explain (y), or so this argument contends” (Sober 1999, 544). In other words, the multiple realizability thesis argues that the higher-level cause *is* explanatory, while the lower-level cause is either *not* explanatory or *less* explanatory than the higher-level alternative.

Before proceeding, it will help to highlight some aspects of Sober’s framework that differ from other treatments of multiple realizability in the literature. In accordance with Sober, I understand realization as a synchronic and symmetric relationship of identity. For example, a particular cigarette may be realized by or identified by chemical A_1 , at a lower-level. Notice how this relationship of realization (or identity) is different from the relationship of causation. As captured in Sober’s figure 2, causation is temporal and asymmetric, as causes both precede and produce their

²As Fodor states, “token physicalism is weaker than reductivism” (Fodor 1974, 100).

³Sober refers to these relationships as “synchronic determination” and “diachronic explanation,” respectively (Sober 1999, 544).




Figure 2: Relationships of realization and causation in Sober's (1999) smoking example.

effects. The same cannot be said of realizers—the synchronic relationship of realizer-to-realized conflicts with the temporal precedence of causes relative to their effects. This is supported by the fact that these realizer-realized relationships fail to meet commonly accepted standards of causation (Woodward 2014, 706-709) and by the fact that we do not claim that the lower-level chemical A_1 causes the cigarette that it instantiates.⁴ While Sober's framework treats realization and causation as distinct types of relationships, other treatments of multiple realizability understand realization as having some sort of causal or functional character (Shapiro 2000; Polger and Shapiro 2016).⁵ While a detailed analysis of each framework is clearly outside the scope of this paper, this clarification can help situate the present analysis in the landscape of current publications on this topic.

How does Sober use the framework represented in figure 3? Sober argues against the multiple realizability thesis by denying that higher-level causes can have an objectively privileged explanatory status over lower-level ones. He frames this discussion in terms of the “amount” of causal detail that might be cited in an explanation, with higher-levels containing less detail and lower-levels containing more. He claims that the amount and level of detail cited in an explanation is dictated by “matter[s] of taste” and scientists's interests (Sober 1999, 551). Attention to scientific practice shows that they sometimes care about more detail, other times they care about less, and there is no fixed rule about some level or amount of detail always being more explanatory than another. In this sense, both x and z should be viewed as explanatorily relevant to y and appealing to each cause supports different explanatory virtues. The higher-level cause provides explanatory breadth, while the lower-level one provides explanatory depth. Neither virtue captures an objective metric for explanatory power. This suggests a kind of flexibility and role for subjective preference in identifying the explanans. The same effect or outcome can have different explanations, which “vary in how detailed they are and in the level of organization described” (Sober 1999, 550).

While Sober denies that higher-level causes are objectively privileged over lower-level ones, he does not argue against such privileging in the opposite direction. Although he can seem to strike a neutral position regarding the explanatory power of higher and lower-level detail, this is not a complete reading of his view. Getting clear on this requires appreciating how he distinguishes explanatory virtues from considerations that make an explanation objectively better than another. He suggests that while different levels of detail can serve different sorts of virtues, none of them latch on to any quality that objectively distinguishes superior explanations from inferior ones. These

⁴Or, that we do not claim that (z) “smoking A_1 ” is a cause of (x) “smoking cigarette P.” Notice that it is challenging to view “x” as causally relevant to “z” because it does not clearly make sense to intervene on “z” with respect to “x”. This is because interventions on z “automatically change” x due to the realization (or supervenience) relation. For further support of this view see (Woodward 2003, 2014).

⁵On these latter accounts, causal or functional facts play a role in distinguishing legitimate instances of multiple realizability from pseudo cases (Polger and Shapiro 2016).

virtues are valued in different contexts without supplying a general, objective metric for gauging explanatory power. In this sense, with respect to explanatory virtues, the level of detail cited can go either way—higher-level detail has virtues that lower-level detail lacks, and vice versa. However, with respect to what makes an explanation “objectively” better than another there is not this type of equality among detail across levels. Sober suggests that there are reasons to think that lower-level details can and do provide objectively superior explanations than higher-level ones. This can be seen in three main points emphasized in his paper. First, he suggests that for any explanatory target, lower-level details can always be included without detracting from an explanation. As he states, lower-level detail is never *unexplanatory* although it might be more than you want to hear. The worst offense committed by this extra detail is that it “explains too much,” while the same cannot be said for higher-level detail. Thus, “physics in principle can explain any singular occurrence that a higher-level science is able to explain...[although]...[t]he level of detail in such physical explanations may be more than many would want to hear, but a genuine explanation is provided nonetheless” (Sober 1999, 561). This reveals a kind of asymmetry in explanatory power: lower-level detail is relevant to any explanatory target, while higher-level detail is rarely (if ever) relevant to lower-level explananda. Second, Sober claims that lower-level details are what really do the “work” in producing higher-level phenomena and that this justifies their privilege or priority in explanations. We see this in his analysis of the smoking example:

“When scientists discover why smoking causes cancer, they are finding out which ingredients in cigarette smoke are carcinogenic. If smoking causes cancer, this is presumably because the micro-configuration of cigarette smoke is *doing the work*. If their [sic] turn out to be several carcinogenic ingredients and different cigarettes contain different ones, this does not make the molecular inquiry explanatorily irrelevant to the question of why people get cancer. The fact that P is multiply realizable does not mean that P’s realizations fail to explain the singular occurrences that P explains. A smoker may not want to hear the gory details, but that does not mean that they are not explanatory” (Sober 1999, 548-549; emphasis added).

Thus, even in cases where higher-level detail is cited, there are always lower-level factors that are really responsible for the regularity or outcome of interest. Both of these points are supported by others in the literature. For example, Waters claims that higher-level detail, while more general, provides “shallow explanations” compared to the “deeper accounts” provided by lower-level detail (Waters 1900, 131). A third reason for the objective explanatory power of lower-level over higher-level detail is that physics has a kind of “causal completeness” that other sciences simply do not have. Sober references Hempel’s work to suggest that completeness is a quality that objectively tracks explanatory power and that it is uniquely provided by physics, as opposed to higher-level sciences. This notion of causal completeness provides an objective measure of explanatory strength, unlike the ways in which higher-level sciences support explanatory virtues or seem “illuminating” for various subjective reasons. This is explained by the fact that “[i]llumination is to some degree in the eye of the beholder; however, the sense in which physics can provide complete explanations is supposed to be perfectly objective” (Sober 1999, 561). Furthermore, “[i]f singular occurrences can be explained by citing their causes, then the causal completeness of physics insures that physics has a variety of explanatory completeness that other sciences do not possess” (Sober 1999, 562). Thus, according to Sober, lower-level detail has more explanatory power than higher-level detail in a way that is untouched by multiple realizability arguments.

3 Multiple realizability in a causal framework. Analyzing the multiple realizability thesis within the framework of causal explanation, as Sober does, is a welcome move. This approach has the potential to capture new challenges that multiple realizability might pose for explanation, how these relate to actual causal scenarios in biology (if at all), and how reductive explanation fares in light of them.⁶ However, there are issues with Sober’s analysis that misrepresent the claims and force of this thesis. Correcting these issues reveals a number of new points. Most notably, it reveals how multiple realizability is related to a particular type of causal complexity that poses problems for reductive explanation, or explanations of higher-level phenomena that appeal to lower-level properties.⁷

Assessing whether some level of detail is causally (and, thus, explanatorily) relevant to an effect requires first specifying two main things: (1) the effect of interest and (2) what is meant by “cause” or “causally” relevant. I will address the causal topic first. In this analysis, I adopt an interventionist account of causation where causes “make a difference” to their effects (Woodward 2003). On this account, to say that C is a cause of E means that some intervention that changes the values of C , and no other variables in background circumstances B , produces changes in the values of E . In other words, causes are factors that operate like “switches” in the sense that they can be manipulated to provide control over their effects. Manipulating these causal factors produces reliable changes in the effects they are related to. This framework can be used to understand the causal relationships represented in Sober’s Figure 2. Within an interventionist framework, we can view variables x , z , and y in the following way. Variable x takes on the values 1 and 0 when a patient smokes and does not smoke, variable z takes on the values 1 and 0 when a carcinogen is present or not, and variable y takes on the values 1 and 0 when lung cancer is present in a patient or absent. Within this set up, to say that x caused y simply means that smoking “makes a difference” to whether one acquires lung cancer or not. The fact that an individual smoked is viewed as the cause of her cancer and if she had not smoked we think that this would have prevented her disease. A similar interpretation makes sense of the causal relationship between z and y , or the lower-level carcinogen and the disease outcome. Variable z is causally relevant to y in the sense that intervening on z —and changing whether the carcinogen is present or absent—provides causal control over whether lung cancer occurs or not. There is nothing mysterious about claiming that the cigarette smoke or the lower-level carcinogen are causally relevant to this instance of lung cancer. However, as we will soon see, this does not mean that these factors are equally explanatorily relevant to all instances of this disease.

Before we examine the relevant explanatory target in this debate, we should fill in the story above and Sober’s Figure 2 with more detail that is present in the multiple realizability case. In this case, the higher-level cause variable x is multiply realized by various lower-level variables, that capture different physical details (z_1, z_2, \dots, z_n) represented in Figure 3. Just like Sober’s original lower-level carcinogen (z), each of these lower-level causes (z_1, z_2, \dots, z_n) are individually sufficient to produce y and, thus, each are causally relevant to some instances of this outcome (y). Figure 3

⁶Although Sober claims that there are legitimate examples of multiple realizability in biology, there is a significant amount of debate in the literature surrounding this view. A significant amount of work has examined how common or uncommon multiple realizability actually is in science, including Bechtel and Mundale (1999), Keeley (2000), Richardson (2008) and Polger and Shapiro (2016).

⁷Notice that this “causal” notion of reductive explanation differs from more traditional accounts, which involve explaining theories by “deriving them or reducing them to other theories” (Wimsatt 1974, 671). Sober relies on a “causal” notion of reductive explanation (Sober 1999, 546), which is compatible with the analysis of causal explanation in this paper.

differs from figure 2 in the sense that it (i) includes the many realizers of the higher-level cause (x), (ii) it displays their causal relation to the higher-level effect (y), and (iii) it represents the higher-level effect (y) as a type-level variable that captures many instances of lung cancer (as opposed to a single instance).⁸ The presence of multiple realizability in this case relates to a unique causal structure. This structure is *causal heterogeneity*—it refers to a situation where distinct instances of the same effect have completely different (or heterogeneous) causes. The causal heterogeneity of y has to do with the multiple realizability of x : the distinct lower-level realizers of x all stand in a causal relationship to instances of y . Causal heterogeneity is very common in biological and biomedical contexts. It is present in cases of genetically heterogeneous phenotypes, which are traits that can be produced by completely different individually sufficient gene variants. I examine these and other cases of causal heterogeneity, in the next section. Notice that figure 3 does not include the lower-level realizers of the effect variable “lung cancer” (similar to Sober’s figure 2). This higher-level effect variable represents a common type of explanatory target in biology and (as I suggest in what follows) one that is a natural focus of the multiple realizability thesis.⁹




Figure 3: Causal heterogeneity.

With these details and the relevant notion of causation specified, we can now consider what the (1) effect of interest or explanatory target should be. In the context of the multiple realizability thesis, one natural explanatory target is a type-level phenomenon in some higher-level science. The target is a higher-level phenomenon because the debate in question is concerned with whether higher-level sciences can be reduced to lower-level ones. Second, the target is type-level because we are considering potential challenges that multiple realizability poses for reduction, and multiple realizability implies a type-level focus. In order to consider the multiple realizations of some higher level property, distinct instances of that higher-level property have to be considered. Multiple

⁸Current literature on this topic often asks: how do we know that a legitimate case of multiple realizability has been identified, as opposed to a pseudo case? Although a sufficient answer to this large question will surely require a separate paper, Sober (1999) and Woodward’s (2014) framework provides a suggestion. In particular, one way to identify legitimate cases of multiple realizability is to fix lower-level properties in order to see if they fix the higher-level entity that they may realize (pseudo realizers should fail this test). For further discussion of this question see (Polger and Shapiro 2016).

⁹How do we know that causal heterogeneity is really different from multiple realizability? As mentioned in section 2.2, this analysis depends on a framework in which realization and causation are distinct relationships. Within this framework, causes precede and produce their effects, while realizers instantiate (and share an identify relation) with the phenomena they realize. If realization and causation are different types of relationships, stating that a property has many realizers is clearly different from stating that this property has many causes.

realization is not visible or apparent at the level of single cases because by definition it involves variation in lower-level realizers across distinct cases. Type-level (or population-wide) explanatory targets are incredibly common in biology and other special sciences. In fact, this can be seen in questions such as “What is the cause of lung cancer?” as this question asks for an explanation of this disease in the human population. Of course, single instances are also of interest in these sciences—we see this in the focus on treating and explaining disease in single patients. However, here the point is that type-level explananda are a focus of the multiple realizability thesis and that they are a representative type of explanandum in actual science. Thus, an issue with Sober’s analysis is that he examines the multiple realizability thesis with respect to explanations of token-level outcomes or single instances, as opposed to considering type-level phenomena. This point is nicely made in Batterman (2018). This ultimately prevents Sober from identifying certain issues that multiple realizability poses for reductive explanation. This is a feature that he retains in his analysis despite the fact that (1) multiple realizability is not evident at the level of single outcomes and that (2) Fodor and Putnam explicitly deny that such cases are a target of their thesis.^{10 11}

Suppose the explanatory target is fixed to a type-level outcome. How does this inform our analysis of what level of causal detail explains and is causally relevant to this target? In the lung cancer example, notice that there is a problem with appealing to a single lower-level carcinogen (such as z_1) in explaining the population-wide occurrence of this disease. The problem is that any single lower-level carcinogen only “makes a difference” to and explains a narrow subset of all cases of the disease. This can be understood in terms of the scope of a factor’s causal control over an effect. In this case, carcinogen z_1 has control of *narrow scope* in the sense that manipulating this factor only controls and explains a fraction of all cases of the effect of interest. In the terminology used above, this can be understood as the carcinogen being a “switch” that only has control over a narrow subset of all cases of the disease. This, of course, is because there are many different lower-level carcinogens that are all capable of causing this disease, so any *single* carcinogen only controls and explains a narrow subset of all cases. This reveals a problem for citing carcinogen z_1 in explaining the type-level occurrence of lung cancer—the problem is that this carcinogen only explains a small fraction of all cases that make up the explanatory target.

Alternatively, notice that the higher-level causal factor “smoking” does not have this problem—this factor does “make a difference” to all (or most) cases of this disease. In other words, this factor has causal control of *broad scope* in the sense that it controls and explains most or all cases of this disease at the population level. Targeting and manipulating the higher-level variable “smoking” controls and explains whether patients in the population get this disease or not. For nearly all cases

¹⁰Fodor and Putnam do not deny the relevance of lower-level physical details to token-level cases, which they consider “token-physicalism” (Fodor 1974, 105), (Putnam 1975, 131-132). They suggest that reduction involves stronger commitments than this type of claim.

¹¹Sober claims that Putnam “discusses the explanation of singular occurrences” (Sober 1999, 552). However, even in the peg-board case, Putnam makes a number of claims that reveal his interest in type-level explananda. These are seen in his claims that increased generality tracks explanatory power and that generality of an explanation means that it “applies to a much more interesting class of systems” (Putnam 1975, 132). In particular: “[t]he fact is that we are much more interested in generalizing to other structures which are rigid and have various geometrical relations, than we are in generalizing to the next peg that has exactly this molecular structure, for the very good reason that there is not going to be a next peg that has exactly this molecular structure. So in terms of real life disciplines, real life ways of slicing up scientific problems, the higher level explanation is far more general, which is why it is explanatory” (Putnam 1975, 132).

in the population, smoking causes this disease and refraining from smoking prevents it *no matter which lower-level carcinogens are involved*. This point makes sense of the fact that biomedical researchers and non-experts appeal to smoking as the cause of lung cancer and explicitly target smoking cessation in efforts to control and prevent this disease. Advertising campaigns and the advice of healthcare professionals are purposefully directed at changing smoking habits, because this allows for control over the population-wide occurrence of this disease. Targeting a single carcinogen in these efforts would be less helpful because this would only offer control over a small percentage of all cases of lung cancer. On the other hand, no matter what different lower-level carcinogens are involved, changes in smoking habits can explain and control the type-level occurrence of this disease.

The scope of a factor's causal control is related to at least two main concepts in the scientific and philosophical literature. First, it is related to the statistical concept of percent variance explained, which has to do with how well an independent variable explains variation in a dependent or effect variable (Utts and Heckard 2007). This notion is frequently represented as a percentage derived from the least squares line and the squared correlation coefficient r^2 . For example, "if a correlation has the value $r = 0.5$, the squared correlation is $r^2 = (0.5)^2 = 0.25$, or 25%... a researcher may write that the explanatory variable *explains* 25% of the variation among observed values of the response variable" (Utts and Heckard 2007, 169). In this context, to say that a variable "explains" some percentage of the variation in another typically means that it "predicts" this variation because it involves correlational as opposed to causal data. However, if there are additional reasons that suggest that the variables stand in a causal relationship—as in the scenario in Figure 3—these notions of "percent variance explained" and "explanation" can take on a causal interpretation. Suppose we want to explain variation in "lung cancer"—we want to know why it varies from present (1) to absent (0) across patients in a particular population. Furthermore, this population contains people who both smoke and do not smoke, where the cigarettes in question contain the causally relevant carcinogens ($z_1, z_2...z_n$) shown in Figure 3. In this population, changes in a lower-level carcinogen, such as z_1 , will explain a smaller percentage of the variance in this target compared to changes in the higher level factor x or changes in smoking habits. There is significant variation in the effect variable "lung cancer" that z_1 does not explain because there are instances of this cancer that are caused by other carcinogens. This is not the case for the causal variable x which relates to smoking habits— x explains a much higher percentage of the variation in the effect because all cases of this disease are caused by changes in x (and all cancer-free cases are caused by its absence). In this population, changes in "smoking" provides a better explanation of lung cancer than changes in " z_1 " because changes in smoking explains a greater percentage of variation in this trait than changes in the particular lower-level carcinogen.

Second, the scope a factor's causal control is also related to Yablo and Woodward's discussions of proportionality or the choice of level in explanation (Yablo 1992; Woodward 2010). These notions involve considerations of proper "fit" between cause and effect variables. In other words, with respect to a specified effect, a cause should provide no more and no less relevant detail. In order to see this, consider Yablo's example of a pigeon that has been trained to peck at a red target (Yablo 1992, 257). As Woodward suggests, one might characterize an explanation of the pigeon's behavior in at least two different ways. Either (a) "the presentation of a scarlet target caused the pigeon to peck" or (b) "the presentation of a red target caused the pigeon to peck" (Woodward 2010). A problem with (a) is that it implies that only the scarlet color causes the pigeon to peck, when this behavior is actually evoked by any shade of red. In this case, too much

irrelevant lower-level detail is being provided, which obscures the factor that is actually doing the causal or explanatory work—namely the color red, not its particular shade. Woodward suggests that the right amount of detail (or “fit”) depends on a causal factor’s ability to provide “accurate information about the conditions under which alternative states of the effect will be realized” and that it provides “*only* such information” (Woodward 2010, 13). As “red” provides accurate information about other situations under which the pigeon would peck and “scarlet” does not, (b) satisfies this criterion better than (a).¹² This case is interestingly similar to the smoking example. Claiming that lung cancer is caused by carcinogen z_1 is misleading because it suggests that this carcinogen in particular, and not others, causes this disease, which is not true. In this case, too much irrelevant lower-level detail has been provided which fails to supply information about other conditions under which the effect of interest would occur. It can suggest that the absence of z_1 will prevent lung cancer, which is incorrect because the presence of z_2 can also cause this disease. What matters for the type-level occurrence of this disease is whether an individual smokes or not, as opposed to the particular carcinogen present in the cigarette. “Smoking” is more proportional to the population-wide incidence of lung cancer than “carcinogen z .” This is because “smoking” provides relevant information (and only such information) about other states of affairs that give rise to this outcome, while “carcinogen z ” does not. This captures how, once an explanatory target is fixed, different causes may have different degrees of “fit” with such a target. If a higher-level cause has better “fit” with an explanatory target than a lower-level one, proportionality captures the objective sense in which the higher-level cause provides a better explanation than the lower-level one.

This suggests a different understanding of the explanatory relevance of higher and lower-level details than Sober claims. According to Sober, lower-level detail enjoys a kind of objective explanatory power that higher-level detail can never provide. If anyone denies the explanatory value of lower-level detail is it probably because such detail is “more than [they] want to hear” or because it explains “too much” (Sober 1999, 547). However, when we look to the smoking and pigeon examples, the problem is not that the lower-level details explain too much, but that they explain too little—they fail to account for all (or most) of the cases of interest. In these situations, it is not that we have heard “more than we want to hear” but that we have not heard enough of what actually matters. We want an explanation of a type-level trait and lower-level details are limited in providing such an explanation. These limitations conflict with Sober’s claims that lower-level details are really doing the “work” and that they provide a type of “causal completeness” that cannot be supplied by higher-level details. In fact, if completeness has to do with accounting for a larger percentage of all cases of interest, the higher-level factors actually provide a more complete explanation than the lower-level ones.

This analysis reveals that multiple realizability does, in fact, pose issues for explanatory reduction. At least one set of these issues can be well understood within a causal framework, in which various lower-level factors are causally heterogeneous for a type-level outcome. In this situation, there are problems for citing lower-level details in explanations and there are objective reasons to privilege appealing to higher-level details over lower-level ones. These reasons have to do with the causal and explanatory relevance of factors to the type-level target and the fact that, in some cases,

¹²Alternatively, if the pigeon did only peck at scarlet targets (and not just any shade of red), (a) would score better along this dimension than (b). This is because (b) would contain too little relevant detail about the conditions under which other states of the effect are realized (it would not explain why the pigeon fails to peck when presented with crimson, but pecks when presented with scarlet.)

the lower-level factors explain fewer instances of this target than higher-level ones. This relates the explanatory power of a cause to the scope of its causal control over a type-level outcome.

One advantage of this analysis is that it provides a straightforward way to understand key concepts employed in some of the earliest arguments for the multiple realizability thesis, such as the notions of generality and scope (Putnam 1975). This analysis differs from earlier attempts to make these concepts clear in two important ways: it formulates them in the context of causal explanation and it defines them with respect to a particular explanatory target (namely type-level phenomena). This makes it clear exactly what “generality” and “scope” refer to and the particular contexts in which they guide determinations of explanatory relevance. Different kinds of explanatory targets are likely to involve different guidelines or criteria for determining which details are relevant and which are not. We can see this with token and type-level targets. Generality, breadth, and scope are highly important for type-level outcomes but less so for explanations of single occurrences. This provides a novel way to understand the operative notions of generality and breadth, which is in terms of the scope of a factor’s causal control over some explanatory target. Furthermore, this indicates a particular way to make sense of Fodor and Putnam’s claims in a causal context. When you are interested in explaining many instances of a higher-level phenomenon, causal details that explain more of these cases provide an objectively better explanation than detail that explains fewer of them.¹³

4 Problems with the disjunctive strategy. It may seem that there is a clear solution to these problems for reductive explanation. Why not explain the higher-level effect by appealing to a disjunctive set of all lower-level causes of y (or realizers of x)? One problem for this approach is that standard views of reduction do not allow for heterogeneous disjunctions (Fodor 1974, 1997). For the project of reducing “ $P \rightarrow Q$ ” the disjunctive set of realizers for each predicate (“ P ” and “ Q ”) will produce disjuncts within any “law” that is formulated at the lower-level.¹⁴ One issue here is that laws are expected to be non-disjunctive and to contain “natural kind terms,” while multiple realization prevents this condition from being met (Fodor 1974). Whether or not these are legitimate issues has been the focus of significant debate (Sober 1999; Butterfield 2011a,b). Much of this debate has centered around explanations that are not clearly causal, but are instead of a deductive-nomological variety or focused on the derivation of higher-level laws from lower-level ones. In keeping with the causal framework of this paper, it will help to consider this disjunctive strategy within the context of causal explanation and, in particular, causal explanations of type-level phenomena.

The question we can consider now is whether in the smoking case, and similar examples, this disjunctive move overcomes the aforementioned issues associated with appealing to lower-level heterogeneous causes. Similar examples are fairly common in the biological sciences. One clear set of examples are cases of genetic heterogeneity, in which distinct instances of a particular phenotype are caused by completely different gene variants (or combinations of them). For example, the disease phenotype retinitis pigmentosa is caused by anywhere from 75-300 different gene mutations and our best scientific evidence suggests that numerous psychiatric disorders, such as schizophrenia,

¹³The feature of “broad scope” should be viewed as one of many criteria that can be used to assess explanatory power—it should be added to lists of other criteria that have been discussed in the literature Ylikoski and Kuorikoski (2010).

¹⁴For example, in reducing the regularity “ $P \rightarrow Q$ ” (from figure 1) we may get a lower-level regularity such as A_1 or A_2 ...or A_n causes B_1 or B_2 ...or B_n .

are marked by similar degrees of genetic heterogeneity (Evans 2017; Uher and Zwicker 2017). So, at the very least, we do find biological cases with this type of causal structure. However, the question remains—in these cases, do appeals to disjunctive sets of causes supply legitimate explanations of type-level phenomena?

The strategy of appealing to disjunctive causes introduces a variety of issues that are helpfully revealed within a causal framework. A first issue with this disjunctive strategy is that it does not appear to accurately capture how explanations work in science. Scientists rarely explain heterogeneous phenotypes by listing off all causes in the disjunctive format suggested. In fact, there are often too many causes to make this a realistic approach. For diseases like retinitis pigmentosa and schizophrenia, this would involve listing off nearly 300 causal factors, which is an absurd requirement to expect valid explanations to meet. If this disjunctive strategy is a viable approach, there is a kind of puzzle in that we do not find it used in scientific practice.

A second issue concerns how to understand the structure of causal explanation when the explanans is a disjunctive set of causes. From the standpoint of an interventionist account, changes in an effect variable are understood as counterfactually dependent on changes in a cause variable. Incorporating the disjunctive strategy into this framework suggests that changes in an effect are counterfactually dependent on changes in a *disjunction*, however it is not clear how exactly to interpret this.¹⁵ In particular, it is not clear what it would mean to intervene on a disjunction (or how one would go about doing this) or what it means to say that an effect counterfactually depends on a disjunction. Such issues apply to all accounts of causation that understand causal relationships in terms of dependency relations. Providing convincing support for the plausibility of the disjunctive strategy requires elaboration on how it works, how it overcomes these challenges, and how it can be reconciled with a specified account of causal explanation.

Third, this disjunctive strategy can be incompatible with the role that unification plays in judgments of explanatory relevance (Lombrozo and Pacer 2017). For example, with respect to explanations of type-level phenomena in biology and biomedicine there is a common view that unified causal explanations are preferable to disunified ones. This is seen in discussions of disease traits, in which there is often a dedicated search for the “shared causal etiology” of a given disease, which refers to some shared causal process that all instances of the disease have in common. The perceived importance of this feature is so strong that it is often used to discount hypothesized diseases from being viewed as “legitimate” and “valid” when they fail to meet it. Clearly, identifying disease traits that meet this shared causal etiology standard is useful from the perspectives of explanation and control. This can satisfy our interest in simple explanations over unnecessarily complex ones (Lombrozo and Pacer 2017). Why tell 300 different causal stories for a disease, when one would suffice? Why therapeutically target 300 different causal factors, when targeting and manipulating one can control all cases of the population-wide disease? When a type-level outcome has a shared causal process this process can be cited in explaining the trait and it can be targeted to potentially prevent, treat, and cure it. These points emphasize how useful it can be to identify shared, single, or unified causes for a given biological trait and how a disjunctive set causes fails to provide this type of utility.¹⁶

A fourth problem with the disjunctive strategy is that it fails to answer an additional question that arises in situations of causal heterogeneity. When it is discovered that a biological trait is

¹⁵This point was first brought to my attention by Jim Woodward.

¹⁶This is suggested by Richardson (2008) who notes that a main issue with disjunction is that it refers to disparate lower-level kinds, as opposed to some higher-level homogeneous property (Richardson 2008, 530).

causally heterogeneous there is often an expectation that a further question should be answered in order to provide a satisfying explanation of the trait. This further question is “why do *different* causes all produce the *same* effect”? There is something puzzling about this situation that conflicts with the intuitive view that the same type of effect should have the same type of cause (Hume 1738, XV).¹⁷ Situations of causal heterogeneity conflict with this intuition and scientists often expect that some explanation of the outcome should furnish an answer to this question or resolve the puzzle. The structure of this scenario is similar to cases of universality discussed by Batterman (2002). In these cases a “universal” behavior is some behavior that is shared across systems with different microstructural details. Examples include microstructurally distinct fluids that exhibit similar features at their critical points and neurons that exhibit the same firing pattern despite having different physical details (Batterman 2002; Ross 2015). Situations of causal heterogeneity are similar to these “universal” behaviors because they capture a shared behavior that is produced by different causal details across different systems. Batterman identifies a particular explanatory why-question that arises in these cases—which he calls a type (ii) why-question—that asks *why* the same behavior is exhibited by systems with different details. Similar to cases of multiple realizability, such a question cannot be answered by citing a particular realizer or cause that is relevant to a single system. This might explain the behavior in a single case, but not why the behavior is exhibited across physically distinct systems. This type (ii) why-question well captures part of the explanatory puzzle in cases of causal heterogeneity. In these cases, there is an interest in knowing why different factors all produce the same effect and citing a disjunctive set of causes fails to answer this question.

These problems show that the disjunctive strategy is not a viable option in situations of multiple realizability and causal heterogeneity for type-level explanatory targets. Thus, in explaining the type-level incidence of lung cancer there are problems for citing single lower-level causes and for citing disjunctive sets of them. As these are clear alternatives to citing a higher-level cause such as “smoking,” and these alternatives face issues, we have more reason to view the higher-level cause as the more compelling explanation.

5 Other scientific examples. This analysis suggests that multiple realizability and causal heterogeneity can pose challenges for reductive explanation. In these cases, higher-level detail can have more explanatory power than lower-level detail. This has been demonstrated by analyzing Sober’s smoking example. In this case, smoking provides an objectively better explanation of the population-wide occurrence of lung cancer than a particular carcinogen. This is because smoking better fulfills the criteria of broad scope and proportionality: smoking explains more instances of lung cancer than any single carcinogen.

While this analysis identifies particular challenges for explanatory reduction it does not support the stronger claims that higher-level explananda are *always* best explained at higher-levels or that higher-level explananda are *always* “autonomous” from lower-levels. While the criteria of broad scope and proportionality indicate which causal detail is most explanatorily relevant to an outcome,

¹⁷As Hume states, “where several different objects produce the same effect, it must be by means of some quality, which we discover to be common amongst them. For as like effects imply like causes, we must always ascribe the causation to the circumstance, wherein we discover the resemblance.” This captures the intuitive view that similar effects should or often have some similar causal story. If this similar causal story is not identified there should be some explanation for how different causes produce the same effect (Hume 1738, XV).

they do not specify the particular “level” at which this detail resides. In cases of causal heterogeneity there are reasons to think that this explanatorily relevant detail can reside at (i) lower-levels, (ii) mid-range levels, or (iii) at the higher-levels of the explanatory target.¹⁸ In order to see this, consider one dissimilarity between Sober’s smoking example and many cases of causal heterogeneity in biology. While many situations of causal heterogeneity involve some higher-level effect (y) and its many lower-level heterogeneous causes ($z_1, z_2 \dots z_n$), they often lack a higher-level variable (x) that all heterogeneous causes instantiate or realize. In other words, these cases of causal heterogeneity look like figure 3, but they lack variable x . This is obvious in examples of genetically heterogeneous disease phenotypes such as retinitis pigmentosa, Parkinson’s disease, and Alzheimer’s disease. These cases involve a higher-level effect (disease trait) and lower-level heterogeneous causes (numerous gene variants), but they lack a clear “higher-level” factor that the genetic causes all realize. These gene variants are lower-level properties that lack a clear higher-level counterpart in the way that lower-level carcinogens relate to the higher-level property “smoking.”

Appealing to such a higher-level, multiply-realized cause was exactly what provided increased explanatory power in the smoking example. If no such cause is present in cases of genetically heterogeneous disease traits, how are such traits explained? One approach involves identifying a causal variable that the heterogeneous factors all causally converge on, and operate through in producing the final effect of interest. Thus, instead of identifying a causal variable that the heterogeneous factors all realize, the goal here is to identify a variable that they all causally converge on. In the biological sciences this is often referred to as a “final common pathway” or “bottle neck” that all causal factors funnel through on their way to producing the single, shared phenotype (Betancur 2011, 62)(Kendler 2012, 7). For example, research suggests that the heterogeneous causes of Parkinson’s disease all converge on the death of dopaminergic neurons, which produces the parkinsonian disease phenotype (Burbulla and Krüger 2011). In Alzheimer’s disease the final common pathway is beta-amyloid accumulation in the brain, and in retinitis pigmentosa it is the death of retinal photoreceptor cells (Gandy 2005, 1121)(Kennan et al. 2005, 108). In explaining the population-wide incidence of these disease traits causal factors in the final common pathway provide more explanatory power than the distinct, heterogeneous causes. This is because such unifying causes better meet the conditions of broad scope control and proportionality, compared to the individual, heterogeneous factors. These unifying causes—located in the final common pathway or convergence point—“make a difference” to all (or most) cases of the disease trait of interest.

This captures a similar explanatory strategy to the smoking scenario, but also an important difference. The shared strategy—found in both the smoking example and these disease examples—involves assessing explanatory power on the basis of the conditions of broad scope control and proportionality. Factors that meet these conditions are often “unifying causes”—they capture some unified causal story that explains all (or most) instances of the type-level outcome. In the smoking example these heterogeneous causes are unified by a higher-level realizer. In the disease examples, the heterogeneous causes are unified by a shared downstream cause that they all converge on. These capture two different ways to unify a heterogeneous set of causal factors and two different ways to explain how different factors all produce the same outcome. With respect to the latter, notice that both of these resolve the puzzle that the disjunctive strategy left open, namely, why different causes all produce the same effect. In the first case, this is because the different heterogeneous factors all realize the same higher-level cause, which produces the same effect. In the second case,

¹⁸I should point out that this conclusion is broadly supported by other work in this literature, such as (Richardson 2008; Keeley 2000; Raerinne 2017).

this is because the different heterogeneous factors all converge on the same causal process, which produces the same effect. Both cases conform to the intuition that the same type of effect should have the same type of cause

However, there are clear differences between these strategies. In the smoking example, the unifying cause is at a higher-level. This leads to the view that higher-level detail can be more explanatorily relevant than lower-level detail. However, in many cases of causal heterogeneity in biology, unifying causes are not always found at higher-levels. In particular, one unique feature of the final common pathway strategy is that these unifying causes can be located at many different levels. These unifying causes can be located anywhere from the (i) lower-level gene variants, to (ii) mid-range cellular processes, to (iii) higher-level behavioral properties. As an example of the first, (i) different gene variants can converge on the production (or misproduction) of the same biochemical product, which captures a lower-level unified cause. Examples of the second situation (ii), of mid-level unifying causes, are seen in the cases of retinitis pigmentosa, Parkinson's disease and Alzheimer's disease. For each of these diseases, lower-level gene variants all causally converge on neuronal and cellular processes, which occupy intermediate biological levels.¹⁹ Finally, an example of (iii) is a situation where these variants all converge on higher-level neural networks, brain circuitry, or cognitive processes, that reflect a higher-level unified cause. Thus, the unifying causal story can be found at different levels for different explananda. The location of the unifying causes captures the location of the "causal action" that explains and ultimately leads to the outcome of interest (Kendler 2013, 1060).

What does this tell us about the "level" of scientific detail that is causally and explanatorily relevant to biological outcomes? Part of what it suggests is a more measured view about the consequences of multiple realizability than claims that the special sciences are "still autonomous after all these years" (Fodor 1997, 149). Factors that are causally and explanatorily relevant to a higher-level outcome can be at the same higher-level, but they do not have to be. Strong claims of autonomy do not make sense of successful and common explanations of higher-level disease traits that appeal to gene variants or cellular processes, which are clearly at lower-levels. Alternatively, all explanations of higher-level phenomena are not reductive—these explananda are sometimes explained by causal factors at the same higher-level or at other mid-range levels that do not collapse down to the lower-levels of biology (or physics). So neither strong higher-level autonomy or lower-level reductive explanation alone captures the nature of biological explanation. Instead, these explanations appeal to detail at various scientific levels. Furthermore, the particular level of detail that biologists cite in their explanations is not determined on the basis of personal preferences, mere "matter[s] of taste," or simply what they are interested in. Instead, the level of detail they appeal to is dictated by the level at which particular causal factors are found—namely, causal factors that meet the standards of broad scope control and proportionality. Once a type-level outcome is selected, it is an objective matter whether a cause or causal process meets these standards or not.

¹⁹This supports a picture, suggested by Schaffner, where biological explanation often involves theories of the "middle range" in the sense that they appeal to phenomena that occupy some mid-level location (Schaffner 2007). In other words, these theories are "not about high level populations evolving in evolutionary time and not about specific DNA sequences or specific enzymes functioning in well defined biochemical pathways, but where at the level of the organelle, the gene as characterized by functional products, the cell, and the organ (Schaffner 2007, 147)." Related ideas are present in Darden and Maull's work on "interfield theories" (Darden and Maull 1977) and Craver's discussion of interfield integration (Craver 2007).

6 Conclusion. This analysis places the multiple realizability thesis within a causal framework and relates it to a common type of causal complexity in biology, namely causal heterogeneity. It is argued that multiple realizability and causal heterogeneity pose problems for reductive explanation. In these cases, lower-level details have less explanatory power than high-level details. This assessment of explanatory power is not made on the basis of personal preference or mere “matter[s] of taste” (Sober 1999, 551). Whether these details are explanatorily relevant or not is determined on the basis of objective considerations—these have to do with whether factors have causal control of broad scope and are proportional to their effects. However, while these identify problems for reductive explanation, this does not mean that they can never be provided or that higher-level explananda are “autonomous” from lower-levels, in some strong sense. In these situations, the level of relevant detail is determined by the level at which some unifying causal story is found.

References

- Batterman, R. W. (2002). *The Devil in The Details*. Asymptotic Reasoning in Explanation, Reduction, and Emergence. Oxford University Press, Oxford.
- Batterman, R. W. (2018). Autonomy of Theories: An Explanatory Problem. *Nous*, 52:858–873.
- Bechtel, W. and Mundale, J. (1999). Multiple Realizability Revisited: Linking Cognitive and Neural States. *Philosophy of Science*, 66:175–207.
- Betancur, C. (2011). Etiological heterogeneity in autism spectrum disorders: More than 100 genetic and genomic disorders and still counting. *Brain Research*, 1380(C):42–77.
- Burbulla, L. F. and Krüger, R. (2011). Journal of the Neurological Sciences. *Journal of the Neurological Sciences*, 306(1-2):1–8.
- Butterfield, J. (2011a). Emergence, Reduction and Supervenience: A Varied Landscape. *Foundations of Physics*, 41(6):920–959.
- Butterfield, J. (2011b). Laws, Causation and Dynamics at Different Levels. *Interface Focus*, 2(1):101–114.
- Craver, C. F. (2007). *Explaining the Brain*. Oxford University Press.
- Darden, L. and Maull, N. (1977). Interfield theories. *Philosophy of Science*, 44:1–23.
- Evans, D. R. (2017). Novel 25 kb deletion of MERTK causes retinitis pigmentosa with severe progression. *Investigative Ophthalmology & Visual Science*, 58:1736–1742.
- Fodor, J. A. (1968). *Psychological Explanation*. MIT Press, Cambridge, MA.
- Fodor, J. A. (1974). Special Sciences (Or: The Disunity of Science as a Working Hypothesis). *Synthese*, 28:97–115.
- Fodor, J. A. (1997). Special Sciences: Still Autonomous After All These Years. *Philosophical Perspectives*, 11:149–163.
- Gandy, S. (2005). The role of cerebral amyloid β accumulation in common forms of Alzheimer disease. *Journal of Clinical Investigation*, 115(5):1121–1129.
- Hume, D. (1738). *A Treatise of Human Nature*. Book 1.
- Keeley, B. L. (2000). Shocking lessons from electric fish: The theory and practice of multiple realization. *Philosophy of Science*, 67:444–465.
- Kendler, K. S. (2012). Levels of explanation in psychiatric and substance use disorders: implications for the development of an etiologically based nosology. *Molecular Psychiatry*, 17(1):11–21.
- Kendler, K. S. (2013). What psychiatric genetics has taught us about the nature of psychiatric illness and what is left to learn. *Molecular Psychiatry*, 18(10):1058–1066.

- Kennan, A., Aherne, A., and Humphries, P. (2005). Light in retinitis pigmentosa. *Trends in Genetics*, 21(2):103–110.
- Lombrozo, T. and Pacer, M. (2017). Ockham’s Razor Cuts to the Root: Simplicity in Causal Explanation. *Journal of Experimental Psychology: General*, 146:1761–1780.
- Polger, T. W. and Shapiro, L. A. (2016). *The multiple realization book*. Oxford University Press, Oxford.
- Putnam, H. (1967). Psychological Predicates. In *Art, Mind, and Religion*. Harvard University Press.
- Putnam, H. (1975). *Philosophy and our Mental Life*, volume 2 of *Mind, Language, and Reality*. Cambridge University Press, Cambridge.
- Raerinne, J. (2017). Explanations of exceptions in biology: corrective asymmetry versus autonomy. *Synthese*, 194(12):5073–5092.
- Richardson, R. C. (2008). Autonomy and Multiple Realization. *Philosophy of Science*, 75(5):526–536.
- Ross, L. N. (2015). Dynamical Models and Explanation in Neuroscience. *Philosophy of Science*, 82(1):32–54.
- Schaffner, K. F. (2007). Theories, models, and equations in systems biology. *Systems Biology*.
- Shapiro, L. (2000). Multiple realizations. *Journal of Philosophy*, 97:635–654.
- Sober, E. (1999). The multiple realizability argument against reduction. *Philosophy of Science*, 66:42–564.
- Uher, R. and Zwicker, A. (2017). Etiology in Psychiatry: Embracing the Reality of Poly-gene-environmental Causation of Mental Illness. *World Psychiatry*, 16(2):121–129.
- Utts, J. M. and Heckard, R. F. (2007). *Mind on Statistics*. Thomson Brooks/Cole, Belmont, CA, 3rd edition.
- Waters, C. K. (1900). Why the anti-reductionist consensus won’t survive: The case of classical Mendelian Genetics. *Philosophy of Science*, 1990:125–139.
- Wimsatt, W. C. (1974). Reductive Explanation: A Functional Account. *Philosophy of Science*, pages 671–710.
- Wimsatt, W. C. (1976). Reductionism, Levels of Organization, and the Mind-Body Problem. In *Brain and consciousness: Scientific and philosophical strategies*, pages 205–267. Plenum Press, New York, NY.
- Woodward, J. (2003). *Making Things Happen*. Oxford University Press, Oxford.
- Woodward, J. (2010). Causation in biology: Stability, specificity, and the choice of levels of explanation. *Biology & Philosophy*, 25(3):287–318.

- Woodward, J. (2014). A functional account of causation; Or, a defense of the legitimacy of causal thinking by reference to the only standard that matters: Usefulness (as opposed to metaphysics or agreement with intuitive judgment). *Philosophy of Science*, 81:691–713.
- Yablo, S. (1992). Mental Causation. *The Philosophical Review*, 101:245–280.
- Ylikoski, P. and Kuorikoski, J. (2010). Dissecting explanatory power. *Philosophical Studies*, 148(2):201–219.