The Explanatory Value of Selecting the Appropriate Scale(s)

Lina Jansson

The Problem

Let me start with a well-worn example of explanations given at different scales. The example is from Putnam (1975). Imagine that we have a rigid board with two holes. One hole is square with sides 1 inch in length. The other hole is circular with a diameter of 1 inch. A solid square peg with sides just less than 1 inch will fit through the square hole, but not through the circular one. Why?

The answer that Putnam (and many others) favours cites the rigidity of the board and the peg together with the geometrical features of the situation: the explanation in terms of macro-structures and properties. We could, in principle, have decided to present the blow-by-blow account of the molecular interactions of the peg and the board: for now, let us call this the explanation in terms of micro-structures and properties. The intuition that Putnam pushes is that even if we did have such a derivation available to us it would not be explanatory (the weaker claim that it is merely a less good explanation will do as the starting point for our discussion).

At least at first glance, this seems like a case where we have gained something valuable by offering the explanation at the macro rather than the micro scale. This observation is often captured by claiming that, for example, generality or abstraction is an explanatory virtue. However, often discussions do not go on to specify how this fits with (or even better, follows from) a broader account of explanation. Why is generality, abstraction, or selecting the appropriate scale(s) explanatorily valuable? The question that I will address in this paper is the following: what kind of explanatory value (if any) comes with selecting the appropriate scale(s)? This question will plunge us into the deep water of debates about emergence versus reduction.

However, even close to shore things are not as simple as they seem. So, let us start in the shallows with: what is scale?

Emergentism versus Reductionism and the Question of Scale

In the very simple case of the square peg and the round hole above, I made use of the most familiar notion of scale: the spatial one. As we move from the macro to the micro we expect to encounter smaller entities and behaviour that takes place at shorter length scales. However, there are plenty of other scales. For example, we might decide to focus on behaviour at higher or lower energies, at longer or shorter periods of time, etc. Given that a spatial scale is not the only scale available, is there anything special about the spatial scale?

Traditional reductionist would answer: yes. The spatial scale is special since it tends to coincide with a necessary condition of parthood for concrete objects. The molecules of the peg are (proper) parts of the peg. Proper parts of concrete entities are (at least in typical

cases of parthood) expected to be spatially smaller than the entities that they compose.[1] The spatial scale turns out to be naturally favoured by classical formulations of explanatory reductionism that make use of the notion of parthood to define what it is to have an explanatory reduction.

As part of their general definition of reduction Oppenheim and Putnam include the criterion that "any observational data explainable by … [the reduced theory] is explainable by … [the reducing theory]" (1958, 5). What makes something the type of reduction that they are interested in, a micro-reduction, is that the reducing theory "deals with the parts of the objects dealt with by…" (1958, 6) the reduced theory. With the assumption that proper parts of concrete entities are spatially smaller than the entities that they compose, calling these reductions *micro-reductions*, seems very appropriate.[2]

According to a traditional explanatory reductionist view, foundational physics can, *in principle*, explain everything that the explanations offered by the sciences outside of foundational physics such as biology, psychology, chemistry or for that matter non-fundamental physical theories such as thermodynamics, etc., can account for.[3] The "in principle" clause requires some spelling out. The idea is not that we actually have to hand replacement explanations from foundational physics nor that our current best candidates for what counts as foundational physics can in principle account for everything explained by all the other branches of science. The "in principle" codifies the expectation that there is a theory that would belong to foundational physics that could, if all the details were known, explain everything currently explained by various different scientific theories in different scientific branches.

Emergentism versus Reductionism and the Question of Explanation

I have only offered a very brief and partial view of the debate between emergentism and reductionism here (see chapters in part 9 of this handbook for a much more detailed discussion). However, it is enough to bring out the role played by the notions of explanation and scale.

> In trying to make emergence intelligible, it is useful to divide the ideas usually associated with the concept into two groups. One group of ideas are manifest in the statement that emergent properties are "novel" and "unpredictable" from knowledge of their lower-level bases, and that they are not "explainable" or "mechanistically reducible" in terms of their underlying properties. The second group of ideas I have in mind comprises the specific emergentist doctrines concerning emergent properties, and, in particular, claims about the causal powers of the emergents. (Kim, 1999, 5)

---

[1] Rueger and McGivern (2010) argue that we should give up on this notion of parthood in the context of understanding explanatory levels. We could, alternatively, give up on thinking of foundational physics as characterised by being a theory of entities that compose the entities of the other sciences. See for example, Ladyman and Ross (2007) for an alternative.

[2] In Kim's (1998) definition we similarly find a focus on parthood.

[3] In many versions primarily from philosophy of science the claim is more specific. The scientific laws, if any, are either part of foundational physics or can be accounted for by foundational physics. In versions primarily from within metaphysics, the focus is instead often on the reduction of macro-properties to micro-properties.

There are several ways that we could approach the explanatory aspect of the emergentist versus reductionist debate. We could focus on case studies and whether we have good reason to take explanatory reductionism to be true in particular cases (for examples of this type of argument, see chapter 36 and 37 in this handbook).[4] Central among the case studies in philosophy of physics is the example of phase transitions; that is, phenomena such as water freezing or the magnetisation of a ferromagnet.[5] These explanations have two interesting features. The first is that the phenomena that we are trying to understand—the behaviours of these systems near criticality—do not appear to be ones that can be understood at a *single* characteristic scale. Second, the explanations offered from the putatively more fundamental theories seem to involve potentially essential appeals to infinite idealisations. Since phase transitions do occur in finite systems, this can be taken to lend support against the claim that these phenomena have been reductively explained by the putatively more fundamental theories.[6]

In addition to arguments for the failure of specific reductive explanations, there are also arguments against the possibility of providing a reductive explanation for some question of interest.

> The challenge of multiple realizability to explanatory reductionism *properly understood,* concerns the ability of the *theory of the heterogeneous micro-realizers* to explain the *robustness* of the common behavior displayed by the systems at macro-scales. That is, the challenge is to explain the *autonomy* of upper-scale common behavior from lower-scale details. However, as we have seen, "disunified" explanations, while certainly telling us a lot about the behavior of individual systems, do not explain the autonomy in question. And, this is true even if we buy into the idea that someday we will have a completed physics—even if we dismiss explanatory difficulties as "merely mathematical" or as involving only "pragmatic" difficulties. (Batterman, 2018, 862)

I will return briefly to this question later. For now, however, I will set aside these much discussed examples. There are two reasons for this. First, I take Knox (2016, 2017) to have convincingly argued that many of the interesting questions about the value of non-foundational explanations are not limited to these cases (even within physics). Her focus is rather on the potential for the selection of appropriate variables to make possible abstractions.[7]

Second, since the approach through case studies is already well-explored elsewhere I would like to present a different route to the questions at stake here. I will seek to clarify what

---

[4] For example, Cartwright (1999), Batterman (2002, 2018), McGivern (2008), Morrison (2012) deny that we have good reason to take explanatory reductionism to be true (although their arguments differ).

[5] These examples are extensively discussed by Batterman (2002, 2013). The discussion of how to understand these cases in relation to explanatory reductionism is very much ongoing. See, for example, Belot (2005), Butterfield (2011a, 2011b), Morrison (2012), Menon and Callender (2013), Reutlinger (2014).

[6] See, for example, Menon and Callender (2013) for a discussion (but not endorsement) of this argument.

[7] The focus on abstraction by being able to reduce the number of variables (and possibly equations) seems to me to be likely to have connections to the more metaphysical approach of Wilson (2010) and the claim that weak emergence can be understood by focusing on degrees of freedom.

kind of explanatory value selecting the appropriate scale(s) *could* be on various broad accounts of explanation.

These two approaches are not, to my mind, competitors. A good account of explanation and of the explanatory value of selecting the appropriate scale(s) will be informed by case studies from science. Similarly, which account of explanation one favours will inform the understanding of the particular case studies.

Three Different Broad Accounts of Explanation

The answer to the question of what the explanatory value of selecting the appropriate scale(s) is will vary depending on the account of explanation that we consider. I will divide accounts of explanation into three different categories (while recognising that the categories are not so sharp as to not admit of borderline cases that are difficult to place): ontic, pragmatic, and epistemic.[8]

On ontic accounts of explanation, explanations are (at least in the primary notion of explanation) in the world.

> Proponents of this [ontic] conception can speak in either of two ways about the relationship between explanations and the world. First one can say that explanations exist in the world. The explanation of some fact is whatever produced it or brought it about . . . Second, the advocate of the ontic interpretation can say that an explanation is something—consisting of sentences or propositions—that reports such facts. It seems to me that either way of putting the ontic conception is acceptable… (Salmon 1989, 86)[9]

On these accounts it makes sense to talk about one aspect of the world explaining some (other) aspect of the world. For example, the hob being turned on explains why the water in the pot is boiling. Typically (as we see in the quote from Salmon), ontic accounts also allow that there is an associated communicative notion of explanation where we display information about which aspect explains what (other) aspect of the world. However, the notion of communication is taken to be derivative on, and secondary to, the idea of explanations as holding in the world.

There are several candidates for what explanations are on an ontic account but some of the most prominent suggestions are that the explanation of some phenomenon should be identified with the cause(s) of that phenomenon (or some subset of the cause(s) of the phenomenon) or, alternatively, with the aspects of the world and the laws of nature that nomologically necessitate the phenomenon to be explained.[10] On this account, there is no great difficulty in discussing explanations as existing independently of us.[11]

---

[8] There are other ways to classify different accounts of explanation. For some related classifications that have informed mine (but are not identical to it) see Jenkins (2008), Marcus (2014), and Bokulich (2016).

[9] For a very similar more recent discussion see Strevens (2008, p. 6).

[10] I am here deviating from Salmon's (1985) three categories.

[11] Of course, subject to the caveat that it is not an explanation about us.

Contrasting with the ontic account, we could put the human aspect of giving, seeking, and having explanations at the centre of our account of explanation. This is what happens on the pragmatic and the epistemic view. Here, the hob being turned on might *cause* the water in the pot to boil and there might be a causal relation in the world linking these two events. However, on these views explanations are not relations in the world.

On pragmatic accounts of explanation, the focus has shifted from taking an explanation to be a relation in the world to the notion of an explanation as an answer to some why-question in some context. Crucially, for pragmatic accounts, whether an explanatory relationship holds depends on the context and the interests of the explanation seeker.

> It is sometimes said that an Omniscient Being would have a complete explanation, whereas these contextual factors only bespeak our limitations due to which we can only grasp one part or aspect of the complete explanation at any given time. But this is a mistake. If the Omniscient Being has no specific interests (legal, medical, economic; or just an interest in optics or thermodynamics rather than chemistry) and does not abstract (so that he never thinks of Caesar's death *qua* multiple stabbing, or *qua* assassination), then no why-questions ever arise for him in any way at all—and he does not have any explanation in the sense that we have explanations. If he does have interests, and does abstract from individual peculiarities in his thinking about the world, then his why-questions are as essentially context-dependent as ours. (van Fraassen, 1980, 130)

On these views, to be told that the hob is turned on might count as an explanation of the boiling of the water in the pot for someone in the right circumstances.

Given the context dependence of the explanatory relation, it is tempting to take pragmatic accounts to be subjective accounts of explanation (in contrast to the objective ontic accounts). Pragmatic accounts of explanations are, however, compatible with a range of views about how dependent on subjective considerations explanations are.

If we take a pragmatic approach that focuses on individual psychological notions of understanding, then we naturally get an account where it is ultimately the interests and background of the individual that determine the success criteria for having an explanation.[12] On these types of pragmatic views what matters is not truth or accurate representation (of at least some aspects of the system of interest) as such. Although, we can make sense of it mattering that the explanation seeker *believes* that the explanation (or some aspect of the explanation) is true or accurately (enough) represents.

Of course, if we expect many features related to achieving understanding to be shared between most humans, we also expect to find a good deal of overlap in terms of the criteria by which explanations should be judged. If, instead of focusing on the notion of individual understanding, we turn to the community level and the idea of intersubjectively shared criteria for explanatory understanding, then we get an account of explanation where the success criteria for explanation are determined by the group at which it is targeted. Here

[12] I take Faye's (2014) account to be of this kind.

the success criteria for explanations are dependent on the explanatory context set by the collective standards and/or interests of the group and not merely on those of the explanation seeker.

Here, it is natural to take it to be the case that "… explanatory relevance is something that is not judged trans-historically (by something like brute number of w-questions) but, rather, is a function of the current state of scientific knowledge" (Bokulich, 2012, 736).

However, pragmatic accounts can also be combined with a general commitment to accurate representation. On such views we still have an account where contextual factors enter into the explanatory relationship itself. One such view is put forward by Potochnik (2010ab).

> Explanation is indelibly context-dependent. … But let me be clear about the extent of this context-dependence. Explanation is only context-dependent at the level of determining which of the many actual causal factors should be included in a particular explanation. If the process under investigation does not conform to some causal pattern, then that pattern cannot explain the outcome of the process. No amount of interest in a pattern can will it into existence. (Potochnik, 2010b, 224-225)

Once we have included an ontic constraint (such as accurate representation of part of the causal history) we have moved far from a pragmatic accounts based around a notion of understanding for an individual. We are now close to the final type of accounts of explanation that I want to consider. This group of accounts differ from pragmatic accounts in that they reject the through and through context-dependence of the explanatory relation. Explanatory relevance can be judged trans-historically on these accounts. However, they differ from the metaphysically focused ontic accounts in taking the role of the agent to be central to explanations (there are no explanations without agents and explanations are not, in the primary notion, relations in the world). On (what I will call) epistemic accounts of explanation context-dependence and interest relativity have a role to play in accounting for our explanatory practices, but they enter into the selection of the explanatory questions that we are interested in. Once we have selected the explanandum of interest (perhaps, including the relevant other options) there is no further role for interest relativity or context to play in determining whether a putative explanans E explains some explanandum; that is, whether an explanatory relation holds between a particular explanandum and a particular explanans is not context dependent or dependent on the interests of particular subjects/communities. A prominent example of the kind of account that I have in mind here is that of Woodward (2003).[13]

> On my analysis, interest relativity enters into what we explain but not into the explanatory relationship itself. What we try to explain depends on our interests, but it does not follow that for a fixed explanandum *M* and fixed explanans *E*, whether *E* explains *M* is itself interest-dependent. Obviously, it is not puzzling and no threat to

[13] Lange (2016) seems like another such view. I also take Strevens' discussion of idealisation, etc., to naturally belong here. I take my own view in Jansson (2015) to be of this kind as well. I also take the unificationist project of Kitcher (1989) and Friedman (1974) to involve relations of explanatory relevance that are not context-dependent.

> the "objectivity" of explanation that the explanans *E* may explain *M* but a different explanans *E*' may be required to account for M'._(Woodward, 2003, 230)

On Woodward's account an explanation must provide accurate information about what would have happened (to the explanandum system) had things been different. While there is scope for interest relativity in selecting the relevant explanandum and the relevant other options that we are considering , there is no further interest relativity in whether or not some putative explanans counts as an explanation of the explanandum in question. Here, some explanans E either explains some explanandum M (once the relevant contrast class, etc. is set) or not, independently of any information about the interests of the agent (or community of agents) seeking an explanation.

I am calling this option epistemic since explanations here are, by definition, something that we (as epistemic agents) seek, receive, and give; they are not primarily simply in the world as in the metaphysically focused ontic account. However, there is no interest relativity or context dependence in the explanatory relation itself, so the account is not pragmatic.

Three Different Accounts of the Value of Selecting the Appropriate Explanatory Scale(s)

With the three broad accounts of explanation on the table (ontic, pragmatic, and epistemic) we can return to the question of what value, if any, comes with selecting the appropriate scale(s) for some putative explanation.

On ontic accounts, explanations are (at least in the primary sense) in the world. Here the correct explanatory scale(s) is naturally thought of as being matched to the scale(s) at which we find the causal or nomological or other modal relations that are taken to be explanatory. If there is explanatory *novel* value at different scales of explanation, on a purely ontic account (setting aside for now that typically a communicative component is recognised too), such value has to be accounted for by the existence of new ontic power (such as new causal or nomological relations) at different scales.

Here we can see the connection between the first and second group of emergentist ideas that Kim (1999) identifies. On ontic accounts of explanation, the first group of ideas demands the existence of the new ontic powers that the second group of ideas focusses on. Of course, this also raises the familiar worries about emergence in terms of causal power that, for example, Kim (1998) raises: how do we accommodate *novel* causal powers by the whole that are not already captured by the parts that the whole supervenes upon?

Within the context of philosophy of physics, the idea that explanatory novelty demands the postulation of a new causal relation leads McGivern (2008, 70) to argue from the independent explanatory role of multi-scale properties for taking very seriously the idea that they are independently causal. This is also the motivation for his consideration of objections to Kim's arguments.

On pragmatic accounts, there are several options for where explanatory value could be found. However, the distinctive option available here but not on epistemic accounts (discussed below) is that the value comes with being sensitive to the interests of the

explanation seeker. This is Sober's response to Putnam's case of the peg that I started this paper with.[14]

> Perhaps the micro-details do not interest *Putnam*, but they may interest *others*, and for perfectly legitimate reasons. Explanations come with different levels of detail. When someone tells you more than you want to hear, this does not mean that what is said fails to be an explanation. (Sober, 1999, 547)

Here, the explanatory novelty involved in selecting the right explanatory scale is simply one to do with matching our interest or disinterest in details at that scale. It does not demand postulating any new causal powers.

The epistemic account introduces a third option that is not available on (purely) ontic accounts. Here the explanatory value of selecting the right scale(s) of explanation can be captured in terms of making available explanatory information that is not available at explanations formulated at competing scales. The explanatory value associated with selecting the correct scale(s) of explanation can derive from our epistemic access to the worldly relations backing explanation. This is neither merely a matter of whether we are interested in the details nor does it demand the introduction of new ontic powers at different scales.

For example, on Woodward's account of explanation (standing in as our representative of epistemic accounts of explanation), the value of selecting the appropriate scale and the appropriate variables is ultimately to be accounted for in terms of the ability to answer what-if-things-had-been-different questions. The value here is neither distinctively metaphysical nor distinctively pragmatic. We can increase our *access* to answers about what-if-things-had-been-different questions both by being correct about what depends on what, by correcting mistaken assumptions about what depends on what, and by seeking presentations that we are well-equipped to grasp.[15]

To see this consider an example where we have a holonomic constraint on motion. Take a bead sliding down a frictionless static wire shaped as a helix (as in picture 1). Let us say that we are interested in the motion of the bead between two points. If we approach this problem on a large enough scale to notice that the helix is of constant radius R and with a uniform pattern of unwinding, then we could use a Lagrangian approach to quickly notice that the motion of the bead is a function of z alone (rather than, say, x, y, z).

---

[14] For Sober (1999) himself this is somewhat more complicated since he does not advocate a pragmatic account of whether some explanans explains some explanandum but merely whether it explains it better or worse.

[15] Ylikoski and Kuorikoski (2010) as well as Woodward (2016) emphasise all of these aspects.

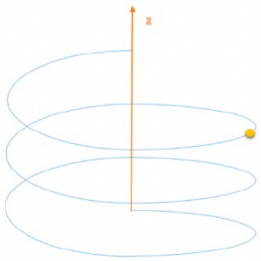

Figure 1

Figure 2

If we instead approach this problem at a smaller scale, this information about the constraints is not always readily available. It is not, of course, that the system does not now satisfy the constraints (it is, after all, the same system as earlier!), but these constraints are not apparent when we approach the problem at the very small scale through, for example, a Newtonian decomposition of forces and calculations of the force from gravity and the normal force (see figure 2).

There are, plausibly, no new ontic powers at the larger scale here. Yet, it is perfectly objective that the motion does not depend on the x and y coordinates and that approaching the problem via a solution in terms of constraints more readily available at the larger than the smaller scale brings this out.
On Woodward's account of causal explanation (standing in as our representative of epistemic accounts of explanation) the value of selecting the right variables lies in avoiding *misrepresenting* the dependences involved by including variables that the explanandum does not in fact depend on.

> [T]he dependence … should be such that (a) it explicitly or implicitly conveys accurate information about the conditions under which alternative states of the effect will be realized *and* (b) it conveys *only* such information—that is, the cause is not characterized in such a way that alternative states of it *fail* to be associated with changes in the effect. (Woodward, 2010, 298)

Here, one of the challenging issues is to allow a principle such as the above, without ending up endorsing as explanatorily optimal explanations that list every possible explanans for the explanandum; typically, such an explanation would look unattractively disjunctive. For example, Franklin-Hall (2016) argues that Woodward's interventionist account does not have the resources to solve this problem without drawing on either pragmatic solutions or solutions in terms more suited to ontic accounts—such as metaphysical naturalness.[16]

Finally, let us return to the challenge from Batterman of explaining the autonomy of the upper-scale from the lower-scale. On an epistemic account this autonomy does not have to

[16] I will leave it open here whether there is an account of naturalness that is neither metaphysical nor pragmatic. However, such appeals to naturalness are also part of, for example, Knox's (2016, pp. 55–57) treatment of the relation between statistical mechanics and thermodynamics.

be understood as a metaphysical autonomy of causal powers. The question here divides into three subquestions. First, do the properties involved in the lower-scale explanation provide the supervenience basis for the properties in the higher-scale explanation? As Potochnik (2010a, 62) has stressed, even if we have supervenience of properties it does not follow that we have supervenience in terms of the properties in *competitor explanations*. The higher-scale *explanation* could involve properties whose supervenience base is not found in the individual lower-scale *explanations* of the phenomenon of interest (although it might be found at the lower-scale). The explanations in terms of a Lagrangian approach with constraints or in terms of a direct Newtonian analysis through knowledge of the normal force would be an example such as this (from the normal force we could not uniquely recover the constraints). There is no particular challenge to global supervenience here, but the properties invoked in the Newtonian explanation do not provide a supervenience base for the properties in the Lagrangian competitor in the particular competitor explanations. Second, assuming that there is an, in principle, explanation of the autonomy of the higher-scale from foundational physics, we have the question of whether this derivation is as good an explanation as the competitor ones. On an epistemic account, this will depend crucially on whether or not we ought to keep our general cognitive abilities fixed when evaluating the "in principle" derivation. Finally, we need to take a stance on questions such as whether the explanatory force stems only from application to nomologically possible scenarios. For example, Weslake (2010) has argued that if logically possible scenarios count, higher-scale explanations can apply to more situations and thus be better explanations than lower-scale ones.[17]

Conclusion

The main goal of this article has been to provide a broad overview of what the value of selecting the appropriate explanatory scale(s) could be on different accounts of explanation. The hope is to have clarified how the account of explanation favoured influences the type of questions that we need to answer in order to address questions about explanatory reduction.

Bibliography

Batterman, R.W., 2002, *The Devil in the Details: Asymptotic Reasoning in Explanation, Reduction, and Emergence*, Oxford Studies in Philosophy of Science, Oxford University Press
Batterman, R.W. , 2013, "The Tyranny of Scales," in R. W. Batterman, ed., *The Oxford Handbook of Philosophy of Physics*, Oxford University Press
Batterman, R.W., 2017 (online), "Autonomy of Theories: An Explanatory Problem," *Nous*, doi:10.1111/nous.12191
Belot, G., 2005, "Whose Devil? Which Details?", *Philosophy of Science*, 72(1)
Bokulich, A, 2008, *Reexamining the Quantum-Classical Relation: Beyond Reductionism and Pluralism*, Cambridge University Press
Bokulich, A., 2012, "Distinguishing Explanatory from Nonexplanatory Fictions", *Philosophy of Science*, 79(5)

[17] Saatsi's chapter on non-causal explanations in this handbook contains examples of cases where it could be argued that non-nomological possibilities should be considered.

Bokulich, A., 2016, “Fiction As a Vehicle for Truth: Moving Beyond the Ontic Conception, *The Monist*, 99(3)
Butterfield, J., 2011a, “‘Emergence, reduction and supervenience: A varied landscape”, *Foundations of Physics* 41(6)
Butterfield, J., 2011b, “Less is different: Emergence and reduction reconciled”, *Foundations of Physics* 41(6)
Cartwright, N., 1999, *The Dappled World: A Study of the Boundaries of Science,* Cambridge University Press
Faye, J., 2014, *The Nature of Scientific Thinking: On Interpretation, Explanation, and Understanding*, Palgrave Macmillan
Franklin-Hall, L. R., 2016, “High-Level Explanation and the Interventionist’s ‘Variables Problem’”, *British Journal for the Philosophy of Science,* 67(2)
Friedman, M., 1974, “Explanation and Scientific Understanding”, *Journal of Philosophy*, 71(1)
Jansson, L., 2015, “Explanatory Asymmetries: Laws of Nature Rehabilitated”, *Journal of Philosophy*, 112(11)
Jenkins, C.S., 2008, “Romeo, René, and the Reasons Why: What Explanation Is”, *Proceedings of the Aristotelian Society*, 108(1)
Kim, J., 1998, *Mind in a Physical World*, MIT Press
Kim, J., 1999, “Making Sense of Emergence”, *Philosophical Studies*, 95(1/2)
Kitcher, P., 1989, “Explanatory Unification and the Causal Structure of the World”, in *Scientific Explanation* edited by P. Kitcher and W. Salmon, University of Minnesota Press
Knox, E., 2016, “Abstraction and its Limits: Finding Space For Novel Explanation”, *Nous*, 50(1)
Knox, E., (forthcoming), “Novel Explanation in the Special Sciences: Lessons from Physics”, *The Proceedings of the Aristotelian Society*
Ladyman, J. and Ross, D., with D. Spurrett and J. Collier, 2007, *Every Thing Must Go: Metaphysics Naturalized*, Oxford University Press
Lange, M., 2016, *Because without Cause: Noncausal Explanations in Science and Mathematics*, Oxford University Press
Marcus, R., 2014, “How Not to Enhance the Indispensability Argument”, *Philosophia Mathematica*, 22(3)
McGivern, P. “Reductive Levels and Multi-Scale Structure”, *Synthese*, 165(1)
Morrison, M., 2012, “Emergent Physics and Micro-Ontology”, *Philosophy of Science*, 79(1)
Menon, T., and Callender, C., 2013,“Turn and Face the Strange . . . Ch-Ch-Changes: Philosophical Questions Raised by Phase Transitions.” in R. W. Batterman, ed., *The Oxford Handbook of Philosophy of Physics*, Oxford University Press
Oppenheim, P. and Putnam, H., 1958, "Unity of Science as a Working Hypothesis", in H. Feigl, G. Maxwell, and M. Scriven (eds*.), Minnesota Studies in the Philosophy of Science*, University of Minnesota Press
Potochnik, A., 2010a, “Levels of Explanation Reconceived”, *Philosophy of Science*, 77(1)
Potochnik, A., 2010b, “Explanatory Independence and Epistemic Interdependence: A Case Study of the Optimality Approach”, *British Journal For the Philosophy of Science*, 61(1)
Putnam, H., 1975, "Philosophy and our Mental Life", in *Mind, Language, and Reality*, Cambridge University Press
Rueger, A. and McGivern, P., 2010, “Hierarchies and levels of reality”, *Synthese*, 176(3)

Reutlinger, A., 2014 (online), “Are causal facts really explanatorily emergent? Ladyman and Ross on higher-level causal facts and renormalization group explanation”, *Synthese*, https://doi.org/10.1007/s11229-014-0530-2
Salmon, W. C., 1985, “Scientific Explanation: Three Basic Conceptions”, PSA 1984, *Philosophy of Science*, 2
Salmon, W. C., 1989/1990, *Four Decades of Scientific Explanation (foreword by Paul Humphreys)*, University of Pittsburgh Press
Sober, E., 1999, “The Multiple Realizability Argument against Reductionism”, *Philosophy of Science*, 66(4)
Strevens, M., 2008, *Depth: An Account of Scientific Explanation*, Harvard University Press
van Fraassen, B., 1980, *The Scientific Image*, Oxford: Clarendon Press
Weslake, B., 2010, “Explanatory Depth”, *Philosophy of Science*, 77(2)
Wilson, J., 2010, “Non-reductive Physicalism and Degrees of Freedom”, *British Journal for the Philosophy of Science*, 61(2)
Woodward, J., 2003, *Making Things Happen: A Theory of Causal Explanation*. Oxford University Press
Woodward, J., 2010, “Causation in Biology: Stability, Specificity and the Choice of Levels of Explanation”, *Biology and Philosophy*, 25(3)
Woodward, J., 2016, “The Problem of Variable Choice”, *Synthese*, 193(4)
Ylikoski, P. and Kuorikoski, J., 2010, “Dissecting Explanatory Power” *Philosophical Studies*, 148(2)