

The explanatory virtue of abstracting away from idiosyncratic and messy detail

Christopher Clarke¹ 

Published online: 19 October 2015
© Springer Science+Business Media Dordrecht 2015

Abstract Some explanations are relatively abstract: they abstract away from the idiosyncratic or messy details of the case in hand. The received wisdom in philosophy is that this is a virtue for any explanation to possess. I argue that the apparent consensus on this point is illusory. When philosophers make this claim, they differ on which of four alternative varieties of abstractness they have in mind. What’s more, for each variety of abstractness there are several alternative reasons to think that the variety of abstractness in question is a virtue. I identify the most promising reasons, and dismiss some others. The paper concludes by relating this discussion to the idea that explanations in biology, psychology and social science cannot be replaced by relatively micro explanations without loss of understanding.

Keywords Explanation · Reduction · Explanatory dispensability · Multiple realizability argument

1 Abstract patterns and understanding

Explanations in the social sciences often employ some concepts that are relatively ‘macro’ compared to the concepts that psychology, for instance, employs. Very roughly speaking, these concepts count as relatively macro because there is a sense in which they are about relatively macroscopic wholes—such as nation states—rather than about their relatively microscopic parts such as individual humans (Clarke Manuscript-a). In turn, explanations in psychology often employ some concepts that are macro relative to those that the physiological sciences employ.

✉ Christopher Clarke
cjc84@cam.ac.uk

¹ Department of History and Philosophy of Science, University of Cambridge, Free School Lane, Cambridge CB2 3RH, UK

And, in turn, explanations in the physiological sciences often employ some concepts that are macro relative to those that the physical sciences employ.

Some of these explanations in physiology, psychology and social science cannot be replaced by any relatively micro explanation without loss of understanding. Or so it seems. The received wisdom in philosophy is that this is because micro explanations include relatively idiosyncratic or ‘messy’ details of the particular case in hand. In contrast, most macro explanations ‘abstract away’ from some of these details. And to abstract away from some of the details of a case, in the right way, is to highlight a ‘pattern’ at the macro level.¹ This is a virtue: explanations that highlight a pattern provide understanding of the case in hand that isn’t provided by explanations that fail to highlight that same pattern. Therefore some macro explanations cannot be replaced by any relatively micro explanation without loss of understanding. (Some would go further, and add that the details included in the micro explanations above are irrelevant, and so such explanations provide no understanding over and above the understanding provided by these macro explanations.² They don’t even complement the understanding provided by these macro explanations.)

This argument from abstractness has played a central role in key debates in the philosophy of biology, mind and social science.³ But, as it stands, the argument cries out for clarification. Firstly, what exactly is it for an explanation to abstract away from the details of another? Or equally, what is it to highlight a pattern? Some authors leave these notions entirely vague (Kitcher 1984; Kincaid 1986; Pereboom and Kornblith 1991; Antony and Levine 1997). Secondly, why is abstractness an explanatory virtue? Thirdly, suppose that understanding is a matter of knowledge; of knowing the causes of a phenomenon for example. How then is it possible that omitting details could improve one’s knowledge and thereby provide extra understanding?

Sections 2, 3 and 4 will untangle four varieties of abstractness that have been knotted together in the literature. For each variety of abstractness, I will examine several *prima facie* reasons to think that the variety of abstractness in question is an explanatory virtue. This clarificatory project will also solve the puzzle of how omitting details could be an explanatory virtue. These sections will also be critical, not just clarificatory. I will dismiss some unhelpful arguments that conclude that abstractness is an explanatory virtue. Indeed, I will argue in Sect. 5 that a supposedly novel variety of abstractness proposed by Haug (2011b) fails to be importantly distinct from the other four varieties of abstractness. Section 6 then shows that these four varieties of abstractness are independent of each other.

¹ See Putnam (1973, 296–297), Garfinkel (1981, 91–96), Kitcher (1984), Marras (1993, 279), Antony (1999, 16) as well as Kincaid (1986, 40–43), Kincaid (1993, 24), Kincaid (1997a) and Potochnik (2010, 69) for talk of ‘capturing’ ‘highlighting’ or ‘bringing out’ patterns. Marchionni (2008) talks of ‘breadth’; and MacDonald (1985, 210) of ‘generality’.

² See footnotes in Sect. 7.

³ See the extensive citations throughout this paper. Indeed, as will be evident from the frequent citation of Fodor’s work, the abstractness thesis is very similar to Fodor’s infamous multiple-realizability thesis (Fodor 1974).

Since my first aim is to explore the explanatory virtue of abstractness itself, I will initially set aside questions about the macro and the micro. Section 7 will then return to this issue and apply the insights from the previous sections to two typical macro versus micro cases.

2 Abstractness as broad jurisdiction of generalizations

This section will focus on the generalizations that an explanation employs. It identifies two respects in which such generalizations can be relatively abstract. But is it a virtue for an explanation to employ generalizations that are abstract in these respects? This section explores a number of prima facie reasons to think so.

Consider the following question: what is the explanatory relationship between mescaline ingestion and hallucination in mammals? Suppose that the threshold dose for hallucination depends upon the species s of the mammal one is considering, its biological gender g , and its body mass m . So one can answer the above question by supplying a function $\tau_1(s, g, m)$ that describes the threshold dose τ_1 of mescaline for each permutation of species, gender and body mass. In other words, one can answer the question by citing the following generalization. (*GI*) “Take any given mammal x and any point in time. The following conditional C holds for any given mescaline level λ above $\tau_1(s, g, m)$, but not for any mescaline level below $\tau_1(s, g, m)$. C : if the mescaline level in mammal x had been λ , then x would have hallucinated.”

Of course, this generalization should be read with an implicit proviso, very roughly: so long as no other causes of hallucination are present. One also wants conditional C to be read here in such a way that generalization *GI* provides understanding of why particular hallucinations occurred, rather than merely providing a means to predict such hallucinations. Hempel and Oppenheim (1948) would say that this conditional is therefore to be read as a claim about the laws of nature: it’s inconsistent with the laws of nature that the individual x receive level of mescaline λ (and several to-be-specified background factors obtain) and yet the individual not hallucinate. Alternatively Lewis (1986) would say that the appropriate interpretation is as a ‘non backtracking’ conditional. And Woodward (2003) would say that it’s as an ‘interventionist’ conditional. The choice here is not relevant to my purposes, and so I leave this question open.

Let’s focus on variable x in the preamble of a generalization such as *GI*, and consider the class of individuals over which x ranges. I will call this class the ‘primary jurisdiction’ of the generalization. The primary jurisdiction of *GI* for example is all actual mammals.⁴

⁴ The generalization ‘All ravens are black’ is logically equivalent to ‘All non-black things aren’t ravens’. So here we have a generalization whose primary jurisdiction is all ravens, but which is logically equivalent to a generalization whose primary jurisdiction is all non-black things. So there’s a sense in which a generalization of the form ‘All Fs are G’ has two sorts of primary jurisdiction. See Sober (1999, footnote 9) for this worry. Note, however, that if the conditionals C in generalizations like *GI* are read as Lewisian or Woodwardian conditionals, then these generalizations will not have the logical form ‘All Fs are G’, and so this problem is averted.

Since $G1$'s primary jurisdiction is all mammals, one might say that $G1$ 'abstracts away' from any particular species. Contrast this with a generalization, $G2$, whose primary jurisdiction is all actual humans. Since $G2$ focuses on a particular species, one might say that $G2$ is less abstract than $G1$. Here's another way of making the same point: $G1$ employs a function $\tau_1(s, g, m)$ that contains three variables; in contrast $G2$ will employ $\tau_2(g, m)$, a function which contains only two of these three variables. On this basis one might say that $G1$ is more abstract than $G2$. In short, breadth of primary jurisdiction is a form of abstractness.

Note that, since it has a broader jurisdiction, generalization $G1$ can be employed in explanations of the hallucinations not just of a human called Eve, but also of a rhesus monkey called Mojo, for example. In contrast, generalization $G2$ can't be employed in explanations of Mojo's hallucinations. This illustrates the trivial thesis that in order for a generalization to provide understanding of each of several cases, its jurisdiction will need to include each of those cases. It follows that (i) there is understanding that a generalization neglects to provide, unless it has a sufficiently broad primary jurisdiction. Indeed (ii) the broader the jurisdiction, the more understanding a generalization can provide, at least up to a point. This is what I will mean when I say that it is a 'virtue' in explanatory contexts for a generalization to have (i) a 'sufficiently' broad primary jurisdiction, and (ii) an 'increasingly' broad primary jurisdiction. (Note that to say that something 'can provide' understanding doesn't entail that it actually does so. Thus, in talking about explanatory virtue in this paper I am pointing to necessary conditions for additional understanding, but not to sufficient conditions.)

This trivial point does not entail, however, that in order to provide understanding of Eve's hallucinations, for example, a generalization will need to include Mojo within its primary jurisdiction, as well as Eve. Thus, whenever a generalization provides understanding of a given case, it is a more controversial thesis to claim that there is additional understanding of that *same case* that the generalization neglects to provide, unless it has a sufficiently broad primary jurisdiction. I will mark the distinction between the trivial thesis and the more controversial one as follows: the trivial thesis is that it is an 'extrinsic' virtue for an explanation to employ some generalizations with a sufficiently broad primary jurisdiction; the controversial thesis is that it is an 'intrinsic' virtue.

Why believe this more controversial thesis? Some philosophers think that an explanation provides understanding only to the extent that it unifies phenomena (Kitcher 1981, 1989). And this principle about understanding suggests that (i) for any given case, there is additional understanding that a generalization neglects to provide, unless it highlights a sufficiently broad pattern. And to highlight a sufficiently broad pattern, in turn, requires the generalization to have a sufficiently broad jurisdiction, such that it includes Mojo for example. This principle also suggests that (ii) the broader the jurisdiction, the broader the pattern can be, and in turn the more understanding a generalization can provide, at least up to a point. Thus it is an intrinsic virtue for an explanation to employ some generalizations with (i) a sufficiently and (ii) an increasingly broad primary jurisdiction. Or so some might argue.

Now, it often goes unnoticed that Fodor (1997, 157–58) proposes an alternative to this unificationist argument.⁵ Fodor claims that good inductive practice is grounded in laws. We observe that all ravens in our sample are black, for example, and we infer that the next raven we observe will also be black; and in some sense, this inference is grounded in the law that all ravens are black. It follows, Fodor claims, that a lawful generalization needs to have a jurisdiction that is sufficiently broad, in order to play this grounding role in induction. That is to say, its primary jurisdiction needs to contain all the things (all ravens, for example) over which the inductive projection is warranted. Next, Fodor implicitly appeals to a principle about understanding: for any given case, there is understanding that a explanation neglects, unless it employs a lawful generalization (Hempel and Oppenheim 1948). But this requires the generalization to have a sufficiently broad primary jurisdiction, Fodor has just argued. Therefore it is an intrinsic virtue for an explanation to employ some generalizations with a sufficiently broad primary jurisdiction. Or so Fodor argues, as I read him.⁶ (Interestingly the analogous point about increasing virtue does not follow, namely: up to a point, the broader the jurisdiction, the more understanding a generalization can provide of a given case.)

Enough about primary jurisdictional breadth. I want to turn now to a second sort of jurisdictional breadth, and thereby a second variety of abstractness. Consider generalization *G3*: “Take any given mammal x and any point in time. The following conditional C holds for any given mescaline level λ between 5000 and 5001 mg, but not for any mescaline level between 0 and 1 mg. C : if the mescaline level in mammal x had been λ , then x would have hallucinated.” Observe that *G1* and *G3* share the same primary jurisdiction. The individual variable x in the preamble of both *G1* and *G3* ranges over the same class of individuals, namely all mammals. Notice *G1*’s predicate variable λ , however. This ranges over a class of state-types, namely all levels of mescaline ingestion. In contrast, the predicate variable λ in generalization *G3* ranges over a much narrower class of state-types, namely all levels of mescaline ingestion between 5000 and 5001 mg, and between 0 and 1 mg.

This draws attention to a second sort of jurisdiction, namely the class of state-types over which predicate variable λ ranges. *G3*, for example, has smaller secondary jurisdiction than *G1*. For another illustration, contrast generalization *G1* with generalization *G4*: “Take any given mammal x and any point in time. The following conditional C holds for all mescaline levels λ_M and all LSD levels λ_L where $\lambda_M + 7.7\lambda_L$ is above τ_1 . But it does not hold for any other levels. C : if the mescaline level in mammal x had been λ_M and the LSD level had been λ_L , then x would have hallucinated.” *G4*’s secondary jurisdiction includes an extra dimension, so to speak, namely state-types in which the individual in question ingests a specified quantity of LSD. *G1* does not include this extra dimension. So *G4*’s secondary jurisdiction counts as broader.

⁵ This is because it is often goes unnoticed that the argument in Fodor (1997) is a considerable the advance over Fodor (1974). For example Jaworski (2002) and Sawyer (2002) treat Fodor 1974 and Fodor 1997 as offering the same argument.

⁶ Note that Fodor talks about explanation implicitly, rather than explicitly. Sober (1999, footnote 17) presents a similar reading to my own, although there are some noteworthy differences.

I distinguish between the primary jurisdiction of a generalization and its secondary jurisdiction because some philosophers are committed to the breadth of this secondary jurisdiction constituting an intrinsic virtue of an explanation, but not the breadth of the primary jurisdiction. For example Woodward and Hitchcock (2003, 190) would say that the more ‘stable’ the functional relationship τ_1 is ‘under interventions’, the more understanding *GI* can provide of any given case. But this just means that (i) for a given case, there is additional understanding that *GI* neglects to provide unless two things hold: firstly, the conditionals *C* in *GI* are true when interpreted in Woodward’s interventionist way and, secondly, the predicate variable λ ranges over a sufficiently broad range of state-types. And indeed (ii) up to a point, the broader the range of state types, the more understanding *GI* can provide of a given case. So, for Woodward, it is an intrinsic virtue for an explanation to employ some generalizations with (i) a sufficiently and (ii) an increasingly broad secondary jurisdiction.

Another way of looking at this is that generalizations with a broad secondary jurisdiction allow one to answer many ‘what if things had been different’ questions, and Woodward (2003) endorses a principle about understanding that says that this ability is an intrinsic virtue of an explanation (Ylikoski and Kuorikoski 2010). For this reason, I think that the distinction between the primary jurisdiction and the secondary jurisdiction of a generalization is worth emphasizing. (Much more on the relationship between a generalization’s primary and secondary jurisdiction in Sect. 6.) Unfortunately the ubiquitous talk in the literature of the ‘scope’ of a generalization tends to obscure this distinction.⁷

In sum, this section has articulated two ways in which one might consider a generalization to be abstract: it has a broad primary jurisdiction, or a broad secondary jurisdiction. This section then drew a novel distinction between intrinsic versus extrinsic virtues. And I noted that many philosophers think that it is an intrinsic virtue for an explanation to employ some sufficiently or perhaps increasingly abstract generalizations. I then identified several general principles about understanding to which one might be tempted to appeal in order to make this point, principles that connect understanding to unification, to laws or to answering what-if questions.

3 Abstractness as logical modesty of the factors cited

Why was the pressure in a given chamber of gas 100 Pa? Answer: because (a) its temperature was 300 k, and its volume was 3 m-cubed; and (b) pressure, volume and temperature are governed by the law that $PV = T$. Section 2 articulated two respects in which the generalizations that an explanation employs, such as *b*, can be relatively abstract. This section will instead articulate a respect in which the particular factors that an explanation includes, such as *a*, can be relatively abstract.

⁷ See Armstrong (1983) and Little (1993) for example for this ambiguous talk of ‘scope’.

And it will explore a number of *prima facie* reasons to think that it is a virtue for an explanation to exclude less abstract factors, or to include more abstract ones.

Firstly, let's ask: why did Eve hallucinate? Modest assertion: she ingested over 1000 mg of mescaline. (This is true.) Bold assertion: she ingested over 1000 mg of mescaline and is a lawyer. (This is also true.) Note that the second assertion logically entails the first one, but not vice versa. That's why the second one counts as logically bolder, and the first as logically more modest. What's more, the modest assertion doesn't contain the detail about Eve's profession. This illustrates an intuitive sense in which logically modest assertions 'abstract away' from logically bolder ones.

Now it seems that the bold assertion provides an incorrect explanation of Eve's hallucinations, unlike the modest assertion. And this is despite the fact that both assertions are true. Somehow, asserting more about the world sometimes makes for an incorrect explanation. What general principle accounts for this curious phenomenon?

Let's make the standard distinction between the semantic content of an assertion versus what the assertion pragmatically conveys. When Sorana asserts for example that the room is hot, she conveys that Sorana wants the window opened. But Sorana does not thereby explicitly assert that this is what she wants. This is a standard case of pragmatics at work. And one can treat the explanation of Eve's hallucination in a similar way, I suggest.

Let's start with the logically modest assertion: Eve ingested over 1000 mg of mescaline. This assertion conveys, I claim, that if Eve had ingested any dose of mescaline over 1000 mg then she would have hallucinated, but if she had ingested any dose under this level then she would not. This generalization is true. Things are different in the case of the logically bold assertion, however: Eve ingested over 1000 mg of mescaline and is a lawyer. One of the things that this assertion conveys is perhaps that (c) if Eve hadn't been a lawyer then she wouldn't have hallucinated, even if she had ingested any given amount of mescaline. Alternatively, one of the things that it conveys is perhaps that (d) if Eve had ingested any given amount of mescaline, but had remained a lawyer, then she would still have hallucinated. But both generalization *c* and *d* are false. So on either reading of what the bold assertion conveys, it conveys a false generalization. (I will call the logically modest and the logically bold assertions 'laconic explanations' to mark the fact that these explanations convey the above generalizations by pragmatic implication, rather than by directly asserting them.)

What's more, *c* and *d* aren't Galilean idealizations, to use Cartwright's terminology. To see this, consider the falsehood that if any two objects were oppositely-charged then they would move closer to each other. This falsehood about oppositely-charged objects is a Galilean idealization in the following respect. The falsehood becomes true if one qualifies it with the proviso 'if all the other factors that cause objects to move were neutralized'. In contrast, *c* and *d* remain false when they are qualified with the proviso 'if all the other factors that cause hallucinations were neutralized'. But I follow Cartwright (2007, chapter 15) in endorsing the following general principle about explanation: Galilean idealizations are more or

less the only falsehoods (if any) that can form part of correct explanations. It follows that the logically bold assertion provides an incorrect explanation.

This illustrates how the pragmatics of explanation-giving allows us to resolve the paradox that asserting a logically bolder truth—one that specifies Eve’s profession—can annihilate the correctness of an explanation. Unless a laconic explanation excludes factors that are too logically bold, the explanation will not be correct, and therefore will not provide any understanding of the case in question. Thus it is an intrinsic virtue for a laconic explanation to exclude sufficiently bold factors. (Note that this reasoning does not establish the analogous claim about increasing virtue, namely: up to a point, the more modest the factors that an explanation excludes, the more understanding it can provide.)

Secondly, let’s consider a slightly modified example: why did Eve hallucinate? Bold assertion: she ingested exactly 2500 mg of mescaline. (This is true.) Modest assertion: she ingested over 1000 mg of mescaline. (This is also true.) Note that the first assertion (that Eve ingested exactly 2500 mg) logically entails the second assertion (that Eve ingested at least 1000 mg), but not vice versa. That’s why the second assertion counts as more logically modest, and thereby as abstracting away from the first assertion. It doesn’t specify the exact dosage that Eve ingested.

Which explanation provides more understanding? Here’s my suggestion. The bold assertion conveys that if Eve were to ingest 2500 mg mescaline then she would hallucinate, but if she ingested no mescaline then she would not. This is ensured by the pragmatics of explanation-giving. But the modest assertion conveys the following: for any λ above 1000 mg, if Eve were to ingest λ mg mescaline then she would hallucinate; but for any λ below 1000 mg, she would not.⁸ Thus the modest assertion conveys a generalization with a broader secondary jurisdiction. So the modest assertion can offer understanding that the bold assertion neglects, according to the Woodwardian principle discussed in Sect. 2. It allows us to answer more ‘what if things had been different’ questions. (And I will endorse Woodward’s principle throughout the rest of this paper as a good measure of understanding, or of at least one aspect thereof.)

This illustrates how (i) for any given case, there is understanding that a laconic explanation neglects, unless it includes some factors that are sufficiently logically modest. Indeed (ii) up to a point, the more logically modest some of these factors are, the more understanding a laconic explanation can provide of the case in question. In other words, it is an intrinsic virtue for a laconic explanation to include some (i) sufficiently and (ii) increasingly modest factors. This time, it’s not a matter of avoiding conveying falsehoods, but a matter of conveying more truths. And it is

⁸ My treatment bears some similarity to Jackson and Pettit’s (1992) treatment of similar cases. One important difference is that my account shows that the logically modest assertion provides both more ‘modally comparative’ information and more ‘modally contrastive’ information—to use their terminology. Jackson and Pettit are mistaken when they claim that logically modest assertions provide more comparative information, and logically bold assertions provide more contrastive information. My treatment also has affinities with that of Marchionni (2008) who also appeals to implicit explanatory contrasts.

not a matter of excluding any relatively bold factors, but of including some relatively modest ones.⁹

At this stage, I should note that there is an alternative way of making the above point that does not rely so heavily on pragmatics. The basic idea is that the logically modest factors cited above ‘made a difference’ to Eve’s hallucinations, but the logically bolder factors did not. One then appeals to the following principle about understanding: an explanation neglects some understanding, unless it cites difference-making factors. However this approach—favoured by Garfinkel (1981, 57–65), Yablo (1992, §8), and Strevens (2008)—uses a very restrictive notion of difference-making. To treat this approach adequately would require another paper in itself; suffice it to say that this restrictive understanding of difference-making is highly controversial.¹⁰

Finally, according to some philosophers there is another reason to regard the inclusion of some relatively modest factors as a virtue. On their view, all other things being equal, an explanation of Eve’s hallucinations is better than an alternative insofar as it can be applied to more cases—for example to Mojo the monkey’s hallucinations.¹¹ Unfortunately it is unclear what these philosophers take the added benefit here to be.¹² Is it merely that the explanation of Eve’s hallucinations can be used as a template for an explanation of Mojo the monkey’s hallucination? Or is it also that the explanation provides understanding of Eve’s hallucination itself that the alternative explanation does not? In other words, do these philosophers take this ability (of being applicable to many cases) to be an extrinsic virtue or instead an intrinsic one?

Some philosophers will think that this ability is an intrinsic virtue because they endorse the unificationist principle about understanding that I discussed in Sect. 2: the more ‘cases’ or ‘systems’ to which an explanation ‘can be applied’, the broader a pattern it can highlight, and so the more understanding it can provide of a given case. Now, talk of ‘applying’ an explanation of one case to another isn’t entirely straightforward, I’d say. But I assume that it requires that an explanation of Eve’s hallucinations, for example, includes some factors that are also present in Mojo’s case, and that also provide understanding of Mojo’s hallucinations. Take for example the factor of ingesting over 1000 mg of mescaline, which provides understanding of both Eve’s hallucinations and of Mojo’s (see Sect. 2). Thus, the more possible cases in which a cited factor is present, the more understanding of a given case can be provided by citing it. In other words, it is an intrinsic virtue for an explanation to include some sufficiently and indeed increasingly modest factors. Or so some would argue.

⁹ To see the difference, consider a hybrid explanation that included both this logically modest factor and this logically bold factor.

¹⁰ See Weatherson (2012), Shaprio and Sober (2012), and Franklin-Hall (Forthcoming, §5) for various problems.

¹¹ See Putnam (1973, 296), Block (1995, §3.3) and Weslake (2010) for this view. Marchionni (2008) and Potochnik (2010) endorses a qualified version of this view too.

¹² See Sober (1999) and Weslake (2010, 291) for illustrations of this ambiguity.

To summarize, this section has examined the particular factors that an explanation includes. And it has articulated a respect—logical modesty—in which these factors can be abstract. This is perhaps the notion of abstractness that Block (1995, §3.3), Weslake (2010) and maybe Putnam (1967, 437) and Kincaid (1990, 63) have in mind.¹³ I examined how a unificationist principle about understanding suggests that it is an intrinsic virtue for any given explanation to include factors that are logically modest. I also examined the special case of laconic explanations. I appealed to pragmatics and to Woodward's principle about understanding to show that it is a intrinsic virtue for a laconic explanation to exclude factors that are sufficiently logically bold; and that it is an intrinsic virtue to include some factors that are sufficiently and indeed increasingly logically modest.

4 Abstractness as syntactic simplicity or cognitive transparency

This section will articulate one final way of defining the abstractness of the parts of an explanation. And I will explore some *prima facie* reasons to think that this final variety of abstractness is a virtue.

Suppose that a chamber of nitrogen gas is connected to a piston. A flame heats the chamber, expanding the nitrogen gas, and thereby moving the piston. Why did the internal energy of the chamber increase? Complex answer: because $(\sqrt{Q} + \sqrt{W})(\sqrt{Q} - \sqrt{W}) > 0$, where Q is the heat that the flame supplies to the chamber, and W is the work done by the chamber on the piston. Simple answer: because the heat supplied by the flame to the chamber was greater than the work done by the chamber on the piston, $Q > W$ in formal terms. Note that $Q > W$ is mathematically equivalent to $(\sqrt{Q} + \sqrt{W})(\sqrt{Q} - \sqrt{W}) > 0$ and so these two answers are logically equivalent. Nevertheless the simple answer has a simpler syntax than the complex answer. Is such syntactic simplicity a virtue of explanations?

Kincaid (1986, 1993, 1997a) and perhaps Marras (1993, 284) seem to argue as follows. (1) It is an intrinsic virtue for an explanation to employ lawful generalizations. This is because such generalizations 'capture patterns'. So the ability of a concept to feature in lawful generalizations is an intrinsic virtue in explanatory contexts.¹⁴ But (2) to feature in lawful generalizations, a concept needs to be sufficiently syntactically simple. Very long disjunctions cannot feature in lawful generalizations, for example. Therefore it is an intrinsic virtue for the concepts that an explanation employs to be sufficiently syntactically simple.

Now, if this argument is to be convincing, it needs to be supplemented with a description of the defining features of lawful generalizations. Otherwise one will be unable to satisfactorily evaluate premise two, the premise that a concept needs to be

¹³ Haug (2011a) may also interpret Fodor (1997) this way, but I'm skeptical of this interpretation.

¹⁴ Kincaid (1986, 40–43) makes an interesting distinction between 'type explanations' and 'token explanations' and says that this argument is focusing on the former. See Marras (1993, 196) for a similar idea.

syntactically simple in order to feature in lawful generalizations. (The appeals in the literature to vague hunches in support of premise two are unhelpful here, I suggest.) Therefore I will consider two putative defining features of lawful generalizations to which Kincaid might appeal. Unfortunately his argument fails on both of these interpretations, I will argue. (The argument also fails for other definitions of lawful generalizations, such as David Lewis' or Marc Lange's (2009, 16), I would also contend.)

The first interpretation is that by a lawful generalization Kincaid just means an explanatory generalization. Take for example the generalization that if any two objects were oppositely-charged then they would move closer to each other. This generalization is explanatory in that it provides understanding of an object's motion (for example according to Woodward's principle about understanding discussed in Sect. 2). However, on this interpretation of Kincaid's argument, notice that premise two comes to mean the following: a concept needs to be syntactically simple in order to feature in explanatory generalizations. But this is basically a re-statement of the conclusion of Kincaid's argument. In other words, on this interpretation, premise one becomes redundant and Kincaid's argument becomes circular. [Indeed this criticism also applies to a similar argument presented in Fodor (1974).]¹⁵

Let's turn therefore to a second interpretation of Kincaid's argument: by a lawful generalization Kincaid means a generalization that has a special sort of evidential status. To illustrate just one sort of special evidential status, take the case in which one discovers a sample of water that freezes at zero degrees centigrade. This sample is a positive instance of the generalization that all water freezes at zero degrees. Now, a popular thought is that, when one discovers any such positive instance, this discovery always lends extra evidential support to the above generalization. Indeed take any second sample of water. A similar thought is that, when one discovers this first positive instance, this discovery always lends extra evidential support to the hypothesis that this second sample will also freeze at zero degrees. To think either of these things is to think that the above generalization has the special evidential status of being 'instance confirmable' (Goodman 1954). In short, my second interpretation of what Kincaid means by a lawful generalization is that he just means a generalization that has a special evidential status, instance confirmability for example.

On this second interpretation, however, I'm strongly inclined to think that premise one of Kincaid's argument is false. That is to say, I deny that a generalization's evidential status makes any contribution to its explanatory virtue.¹⁶ The ability of a generalization to help one understand a given phenomenon does not

¹⁵ It is fair to interpret Fodor (1974), as Kincaid does, as making a similar claim: syntactic simplicity is required for a concept to feature in explanatory generalizations. But Fodor confesses that his argument for this claim is somewhat circular (102). This is in contrast to Fodor's more developed (1997) treatment, which I discussed in Sect. 2.

¹⁶ For one thing, see Sober (1988) for a compelling argument that the importance of instance confirmability has been overstated. I should also note, however, that my contention here leaves open the question of whether evidential status can ever serve as a rough *indicator* of explanatory virtue. Antony (1999), for example, would say that instance confirmability indicates that a generalization refers to natural kinds, and thereby indicates its explanatory virtue.

depend, I'd say, on how one came to know the generalization, or upon how one might have come to know it. Indeed in another context Kincaid agrees; he is adamant that these two things are independent (Kincaid 1996, 94). So Kincaid's argument fails on this second interpretation, and by his own lights. (To those who remain unconvinced, let me offer another reason to think that Kincaid's argument fails: on the present interpretation, the second premise of the argument comes to mean: a concept needs to be syntactically simple, if it is to be part of a generalization with a special evidential status. This premise is also dubious, I'd say. This is because two logically equivalent propositions always enjoy the same degree of evidential support, as is generally acknowledged. It follows that if one of these two generalizations enjoys a special evidential status, then the other does too; even if the first is much more syntactically simple than the second.)

Let's take stock. We are examining Kincaid's hunch that syntactically simple explanations can provide understanding than syntactically complex ones cannot. For example, the hunch is that the simple $Q > W$ answer provides understanding that the complex $(\sqrt{Q} + \sqrt{W})(\sqrt{Q} - \sqrt{W}) > 0$ answer does not, when it comes to explaining why the internal energy of the chamber increased. But, so far, this remains a bare-faced contention. We do not yet have a satisfying argument that justifies this contention or that accounts for it. It's time to put this right.

Note that the syntactical simplicity of the $Q > W$ answer makes its logical implications more cognitively transparent than those of the more complex $(\sqrt{Q} + \sqrt{W})(\sqrt{Q} - \sqrt{W}) > 0$ answer. For example, the simple answer makes it transparent that the energy of the chamber would increase if $Q = 3$ and $W = 2$. The complex answer does not make this relationship cognitively transparent; it obscures it from us, or at least from any cognitively normal person. But knowing this relationship between Q and W and the energy increase provides understanding of the energy increase. (On Woodward's principle, for example, this is because this relationship tells us what would happen if Q or W were different.) Therefore the simple answer makes more transparent the knowledge that we require to explain the energy increase, unlike the complex answer. And so the syntactically simple answer provides more understanding of the energy increase than the syntactically complex one. (And this is despite the fact that both explanations are logically equivalent.)

What's more, one might feel that there is a sense in which the syntactically simple and cognitively transparent answer 'abstracts away from the messy details' of the complex answer. Thus the above discussion illustrates how abstractness qua cognitive transparency is an intrinsic virtue: the more cognitive transparency an explanation provides, the more understanding an explanation can offer of a given case. This also illustrates how abstractness qua syntactic simplicity is an intrinsic virtue—for cognitively normal people—because syntactic simplicity is required for cognitive transparency. So Kincaid's hunch is vindicated, even though his argument is not.

I hope that this example also makes clear that, for cognitively perfect beings, syntactic simplicity is not an explanatory virtue. Imagine, for example, a mathematician for whom mathematical reasoning is so effortless that the $Q > W$ explanation and the $(\sqrt{Q} + \sqrt{W})(\sqrt{Q} - \sqrt{W}) > 0$ explanation are equally

cognitively transparent. In this case we have two entirely synonymous and equally cognitively transparent explanations. Now, understanding depends upon precisely two things, I contend: what the explanation in question says about the world, and how the explanation is cognitively processed. But, in this case, the two explanations say logically equivalent things about the world, and are cognitively processed in the same way. So neither explanation provides understanding that the other does not provide. This illustrates how cognitively perfect beings can have complete understanding without syntactic simplicity. For them it is not an explanatory virtue.

5 Abstraction and causal profiles

I now move on to consider a final variety of abstractness. This variety is formulated by Haug (2011a, b) in an ingenious, if rather tricky, body of work. What follows is my best attempt at a simple and fair exegesis. Haug wants his variety of abstractness to be importantly distinct from the other varieties discussed so far. I will conclude, however, that this fails to be the case.

I will take the ingestion of morphine as an illustration. (Even though the example is my own, I will pretend it is Haug's to avoid cumbersome locutions such as 'Haug would say' and the like.) Now, any ingestion of morphine has the power (i) to slow down the heart and lungs, and in extreme amounts to induce organ failure; and it has the power (ii) to relieve pain and to inhibit pain aversion behaviours.¹⁷ In other words morphine is both (i) a vaso-suppressant and (ii) an analgesic. For simplicity of illustration let's pretend that (i) and (ii) describe precisely those causal powers that are shared by all morphine ingestion events, as dictated by the laws of nature. These are all and only the causal powers 'nominally associated' with the property of morphine ingestion, as I will put it. They constitute the 'full nomic profile' of the property of morphine ingestion.

Again to keep the example simple, let's also pretend that morphine is the only nominally possible vaso-suppressant.¹⁸ So every ingestion of a vaso-suppressant is the very same event as an ingestion of morphine, and vice versa; or so one might infer.¹⁹ Therefore, the property of morphine ingestion and the property of vaso-suppressant ingestion are, by my definition, nominally associated with exactly the same powers. Each of their full nomic profiles is described by (i) and (ii), to be specific (Haug 2011b, 253, 257).

Now imagine that Juliet drinks a potion, her heart stops, and she dies. Why did Juliet die? Answer *MO*: she ingested morphine. (This is true.) Answer *VS*: she ingested a vaso-suppressant. (This is also true.) Haug says that *VS* adds to the

¹⁷ No worries if you think that it's ingestees not ingestions that possess the powers.

¹⁸ As Haug (2011b, §5) urges, however, the following reasoning would still apply if one drops this pretense. Simply replace 'morphine' with a very long disjunction of all the possible vaso-suppressants.

¹⁹ Although Haug (2011b, 253, 257) accepts this, one might dispute this. One might prefer instead to say that the ingestion of morphine is a distinct event from the ingestion of a vaso-suppressant. It's just that the two events are necessarily concurrent. But see Clarke (Manuscript-b) for an argument that my conclusion follows anyway: necessarily concurrent events have exactly the same causal powers.

understanding provided by *MO*. This is because *VS* abstracts away from *MO*, Haug claims.

In what respect does *VS* abstract away from *MO* according to Haug? Haug (2011b, 259) claims that there is an intuitive and important sense in which vaso-suppressant ingestion isn't associated with the power to relieve pain. Let's say this property is not 'Haug-associated' with the power to relieve pain. In contrast, the property of morphine ingestion is Haug-associated with the power to relieve pain, he claims. But Haug then contends that the power to relieve pain is irrelevant to the explanation of Juliet's death. (Instead it's the potion's power to slow down the heart and lungs that is relevant.) So there's a sense in which the *MO* explanation appeals to causal powers that are irrelevant to explaining Juliet's death. In contrast, the *VS* explanation 'abstracts away' from the causal powers that are irrelevant to the explanation. Therefore the *VS* explanation provides understanding of Juliet's death that the *MO* explanation does not provide, Haug argues (259). And so the above variety of abstractness is a virtue of explanations, Haug argues: without it, some understanding is neglected.

But what does Haug association amount to, one might wonder? It is evident from my exposition above that Haug is committed to the following claim: a causal power (relieving pain) can be nomically associated with a property (vaso-suppressant ingestion) without being Haug-associated with that property.²⁰ So Haug is working with a conception of association that is not nomic association. Unfortunately Haug leaves it very much open what it is for a property to be Haug-associated with a causal power.

Putting this serious interpretative worry aside, I'm happy to endorse Haug's conclusion that *VS* provides some understanding that *MO* does not. Here's my own account of the difference. *VS* makes the potion's power to slow down the heart cognitively salient. And this influences the pragmatics of the *VS* explanation: *VS* hints that if Juliet's heartbeat hadn't been slowed down, then she would not have died—for example if Juliet had been given a dose of naloxone immediately after drinking the potion.²¹ And this true proposition provides understanding of Juliet's death. (For example, it answers an important 'what if things had been different question'.) *MO* in contrast does not do this. Instead *MO* makes the potion's power to relieve pain cognitively salient. So if anything, *MO* hints that if Juliet had been in pain, then she would not have died—for example if her dorsal posterior insula had been directly stimulated. And this proposition is false. So *VS* and *MO* hint at different things. And, as a result, *VS* makes more cognitively salient the knowledge needed to provide understanding of the case, in contrast to *MO* which perhaps makes a false proposition more cognitively salient.

To underline my point, let's consider two further explanations of why Juliet died. Answer *VS*⁺: she ingested a vaso-suppressant, and such ingestions are able to slow down the heart and lungs. Answer *MO*⁺: she ingested morphine, and such ingestions

²⁰ This is tantamount to Haug's (2011b, 253) rejection of what he calls the Absolute Closure principle.

²¹ Perhaps *VS* doesn't quite pragmatically imply this, but *VS* at least makes this proposition more cognitively salient.

are able to slow down the heart and lungs. I hope that the reader shares my hunch that MO^+ provides at least as much understanding of Juliet's death as the VS^+ explanation. (According to the account I've just given, this is because both VS^+ and MO^+ highlight the power of the potion to slow down the heart and lungs. Thus both explanations provide the same understanding, in that they make cognitively salient the fact that if Juliet's heartbeat hadn't been slowed down, then she would not have died.)

I conclude that the only difference between MO on the one hand, and VS and VS^+ and MO^+ on the other, is that MO doesn't make the power of the ingested potion to slow down the heart and lungs cognitively salient. It follows that Haug abstractness (whatever it is) is only an explanatory virtue insofar as it provides this sort of cognitive saliency. Therefore Haug abstractness fails to be importantly distinct from abstractness as cognitive transparency that I discussed in Sect. 4. Instead it's just a special case of the latter.

6 The independence of the varieties of abstractness

Having articulated four respects in which part of an explanation can be abstract, I now want to note that these varieties of abstractness are independent of each other in the following sense. One can make part of an explanation less/more abstract in one of these four respects without making it less/more abstract in the other three respects, in principle if not in practice. (I talk of parts of explanations because explanations are complex: an explanation can have one part that is more abstract than another explanation, and a second part that is less abstract.)

Consider the following explanation of why Eve hallucinated: (a) Eve is a female human weighing 70 kg, (b) Eve ingested over 1000 mg of mescaline; (c) for any mammal and any mescaline level λ above $\tau_1(s, g, m)$, if the mescaline level in mammal x had been λ , then x would have hallucinated; where (d) $\tau_1(\text{human, female, } 70) = 1000$ mg. Note that one can narrow the primary jurisdiction of generalization c by replacing 'for any mammal' with 'for any human'. One can narrow the secondary jurisdiction of generalization c by replacing 'any mescaline level λ above $\tau_1(s, g, m)$ ' with 'any mescaline level λ above $\tau_1(s, g, m) + 4000$ mg'. One can make some of the cited factors less logically modest by replacing 'ingested over 1000 mg' with 'ingested precisely 1500 mg' in b . And one can make b less cognitively transparent by rewriting 'over 1000 mg' with the mathematically equivalent description:

$$> 10 \left(\sum_{i=1}^4 i \right)^2$$

Note that each of these changes is independent of the other. For example, the above change to the cognitive transparency of b doesn't change its logical modesty. Nor, for example, does narrowing the breadth of the primary jurisdiction of c (to only humans) change its breadth of secondary jurisdiction (all levels of mescaline ingestion above 1000 mg).

I want to add two caveats to my independence claim here. Firstly, my independence claim does not deny that, in practice, parts of an explanation that are abstract in one respect may tend to also be abstract in other respects. Imagine that one conjoined ‘Eve ingested over 1000 mg of mescaline’ in b with ‘Eve is a lawyer’ to form ‘Eve ingested over 1000 mg of mescaline and is a lawyer’. This change makes b less abstract in two respects. It now employs logically bolder factors, and it is now more syntactically complex. So, conversely, there will be cases in which the most natural way of making part of an explanation more abstract—removing ‘Eve is a lawyer’—makes it more abstract in at least two respects.

Secondly, my independence claim does not deny that there is sometimes an interesting evidential relationship between a generalization’s primary and secondary jurisdictions. For illustration, take any iron surface that is actually exposed to oxygen, and call it x_1 . And take the state-type λ_2 of being damp. Suppose one knows that any such surface would rust if it were damp. And now take a second iron surface that is actually damp, and call it x_2 . And take the state-type λ_1 of being exposed to oxygen. One might reasonably infer from our knowledge about the first surface that this second surface would rust if it were exposed to oxygen.

It is interesting to note, however, that such inferences won’t always be reasonable. Example one: suppose one knows that any given person who actually lives in a region of the Himalayas (x_1) would donate money to charitable causes if they were (λ_2) to be given a million dollars. It might not be reasonable to infer, however, that any actual millionaire (x_2) would donate money to charitable causes if they were (λ_1) to live in this region of the Himalayas. Example two: any given person who actually has an XY chromosome (x_1) would have bad lungs if they were (λ_2) to smoke cigarettes. This is a straightforward supposition. But it’s not at all straightforward to say of any actual smoker (x_2) that they would have bad lungs if they were (λ_1) to have an XY chromosome. This is because it’s not clear that it makes sense to entertain the possibility of an individual changing their chromosomal make-up from XX to XY.

In summary, one can make part of an explanation less/more abstract in one of the four respects without making it less/more abstract in the other three respects. However, there will be cases in which the most natural way of making part of an explanation less/more abstract will indeed make it less/more abstract in two or more respects.

7 How abstract are typical macro explanations?

Having articulated four ways in which part of an explanation can abstract away from another, I want to apply the lessons learned to two typical cases of a macro explanation versus a micro explanation. I must leave it as a task for another day, however, to consider how this bears on macro versus micro cases that are dissimilar to the following two cases (Potochnik 2010; Clarke Manuscript-a). (This task is an important one: to show in this manner that some macro explanations are not replaceable by any micro explanation without loss of understanding, one would have to show that some macro explanations abstract away from *all* micro

explanations of the same phenomenon. I also note that this task is complicated by the ambiguity surrounding what counts as a macro concept or micro concept, and therefore what counts as a macro explanation or a micro explanation.)

Case one In 1990 over one million people carry the herpes virus. Imagine an epidemiologist who has the resources to consider one million carriers in 1990 individually. For each of the million carriers she traces the idiosyncratic sequence of infections that led to the carrier's being infected. In Patrick's case for example she traces the sequence of herpes infection from Vinesh to Alex to Caitlin to Patrick. (And to explain why Caitlin infected Patrick, for instance, she employs the generalization that for all individuals, if they were to have unprotected sex with someone carrying the virus, the chances of being infected is high.) Thus this epidemiologist identifies one million such causal sequences. And these sequences, taken together, were sufficient in the circumstances for there to be at least one million carriers. The epidemiologist claims that this explains why there were at least one million herpes carriers in 1990. I am going to assume that this explanation is correct; if only for illustrative purposes.²²

Contrast this with a second epidemiologist who correctly explains why over one million people carried herpes in 1990. The epidemiologist notes that the society in question is a conservative society: condoms are not freely available, and their use is frowned upon. The epidemiologist also notes that the network of sexual relationships in the society has a 'scale free structure'. This epidemiologist then provides a general equation that links the incidence of the virus at any given point in time, to the incidence of the virus a year later, given the scale-free structure of the society and given the society's conservativeness. One instance of this generalization is the following: if there were over 900,000 carriers in 1989 (which there were) then there would be over one million carriers in 1990. Thus the second epidemiologist explains why there were over one million carriers in 1990.

Note that this second explanation employs some concepts that, one is inclined to say, are about whole populations, such as a society's being conservative or scale-free. The first explanation only uses concepts that are about individual people, in contrast. Thus most philosophers would count the second explanation as macro relative to the first explanation. So let's ask whether any parts of this macro explanation abstract away from any parts of this micro explanation. Sects. 2 to 4 have articulated four varieties of abstractness, so we can make this question more precise as follows.

(I) Take the individual variable x in the preamble of the generalization that the macro epidemiologist employs. Does x range over a broader class of things than the generalization that the micro epidemiologist employs? No. In fact, the two classes are disjoint, strictly speaking. The micro epidemiologist employs a generalization

²² Many would argue that this epidemiologist has not identified what caused this phenomenon, but rather what caused each of the mereological parts of this phenomenon, as it were. See Putnam (1973, 296–298) and Garfinkel (1981) for advocates of this extreme skepticism. Jackson and Pettit (1992), Kincaid (1997a), Sober (1999) and Marchionni (2008) repudiate it. See Owens (1989) for an excellent discussion of the general issues involved. See Kitcher (1984), MacDonald (1992, 86, 90–92) and Haug (2011a, 1150) for the claim that the present explanation includes irrelevant details.

that generalizes over the class of all individual humans; the macro epidemiologist employs a generalization that generalizes over the class of all populations of humans.

(II) Similarly, does the predicate variable λ in the preamble of the macro generalization range over a broader class of state-types? No. In fact, the two classes are again disjoint, strictly speaking. The micro epidemiologist employs a generalization that generalizes over the following state types: ‘unprotected sex’ versus ‘protected sex’, and ‘sexual partner already infected’ versus ‘sexual partner not already infected’. In contrast, the macro epidemiologist employs a generalization that generalizes over the following state types: ‘conservative’ versus ‘not conservative’, and ‘scale free’ versus ‘not scale free’. (Of course, the macro epidemiologist’s generalization answers what-if questions additional to those answered by the micro epidemiologist’s generalization, and thus it provides additional understanding, on Woodward’s principle from Sect. 2. But this is not because it is more abstract. It isn’t, I’ve just shown.)

(III) Are some of the particular factors cited in the macro explanation more logically modest than the micro explanation? Yes. Consider the thesis that 900,000 particular (named) individuals carry herpes. Obviously this thesis entails that at least 900,000 individuals carry herpes. But the reverse does not hold. The fact that at least 900,000 unspecified individuals carry herpes does not entail that it’s the individuals named above who carry it. So the factor cited in the macro explanation “there are at least 900,000 carriers in 1989” is more logically modest than the factor cited in the micro explanation, which identifies 900,000 carriers in 1989 by name. (However, since the micro epidemiologist’s explanation is not a laconic one—which conveys explanatory generalizations by pragmatic implication—the significance of this point is questionable; see Sect. 3.)

(IV) Do the simple syntactic features of the macro explanation make it more cognitively transparent than the micro explanation? Yes. The micro explanation is very complex, describing one million particular causal sequences of herpes transmission.

In sum, in this typical pair of explanations, the macro explanation abstracts away from the micro explanation in two respects, but as far as the other respects are concerned, neither explanation abstracts away from the other.

Case two On 9th April 2015 the cover story of Time Magazine was entitled ‘Black Lives Matter.’ A micro explanation of this event is that five days previously Michael Slager (a police officer) killed Walter Scott (an African-American citizen) by shooting him in the back eight times. A macro explanation of this event is that police forces in the United States (continue to) use excessive force against their citizens, in particular against African-Americans. Many would count the latter as a relatively macro explanation because it refers to whole police forces and whole populations, rather than to two individuals.

Both explanations are equally syntactically simple, more or less. And neither explanation appeals to any generalization, and so both have the same (trivial) primary and secondary jurisdiction. Where the explanations do differ, however, is the logical modesty of the factors that they cite. Take the fact that a police officer called Slager shot an African-American citizen called Scott eight times in the back;

this entails that police forces in the United States sometimes use excessive force against African-Americans. But the reverse entailment doesn't hold. And so the macro explanation abstracts away from the micro explanation, in the sense that some of the factors that it employs are more logically modest.

8 Implications

The notion of an abstract explanation plays a central role in the philosophy of biology, mind and social science. But many authors leave the notion of abstractness entirely vague (Kitcher 1984; Kincaid 1986; Pereboom and Kornblith 1991; Antony and Levine 1997). This paper has articulated four ways of defining the abstractness of the parts of an explanation. This follows in the footsteps of Sober (1999) and Weslake (2010), who attempt something similar. By drawing together disparate work across the philosophy of mind, biology, and social science, this paper offers a more comprehensive taxonomy. Its taxonomy is also an advance on Sober and Weslake's in that it draws a number of helpful distinctions: between a generalization's primary and secondary jurisdiction; between extrinsic and intrinsic explanatory virtues; between the virtue of sufficient abstractness and of increasing abstractness; and indeed between the virtue of including more abstract factors and the virtue of excluding less abstract ones. It also distinguished the question of what is virtuous for laconic explanations, what is virtuous for explanations given by cognitively normal humans, and what is virtuous for all explanations in general. It also separated the task of characterizing abstractness, the task of appealing to general principles about understanding in order to evaluate its explanatory virtue, and the application of these results to the micro–macro case.

Thus this paper has shown that the apparent consensus in the philosophical literature on macro explanations is illusory. Philosophers differ on how they understand abstractness, and for each variety of abstractness there are several alternative reasons to think that the variety of abstractness in question is an explanatory virtue (Sects. 2, 3, 4). Having untangled these varieties of abstractness, and these several reasons to think abstractness a virtue, one was able to solve the puzzle of how omitting details could be an explanatory virtue (Sects. 2, 3). My aims in this paper, however, have not only been clarificatory but also critical. I have dismissed some unhelpful arguments from Fodor, Marras and Kincaid to the effect that abstractness *qua* syntactic simplicity is an explanatory virtue (Sect. 4). And Sect. 5 argued that Haug's supposedly novel variety of abstractness fails to be importantly distinct from abstraction as cognitive transparency.

Acknowledgments I am indebted to Christopher Cowie, Tim Lewens and Nick Shea for their helpful comments on an early draft of the manuscript. This work has received funding from the European Research Council under the European Union's Seventh Framework Programme (FP7/2007-2013)/ERC Grant agreement no 284123.

References

- Antony, L. M. (1999). Multiple realizability, projectability, and the reality of mental properties. *Philosophical Topics*, 26, 1–24.
- Antony, L. M., & Levine, J. (1997). Reduction with autonomy. *Philosophical Perspectives*, 11, 83–105.
- Armstrong, D. M. (1983). *What is a law of nature?*. Cambridge: Cambridge University Press.
- Block, N. (1995). The mind as the software of the brain. In D. N. Osherson, L. Gleitman, S. M. Kosslyn, S. Smith, & S. Sternberg (Eds.), *An Invitation to cognitive science* (pp. 170–185). Cambridge: MIT Press.
- Cartwright, N. (2007). *Hunting causes and using them*. Cambridge: Cambridge University Press.
- Clarke, C. Manuscript A. How to define higher level explanations. Unpublished Manuscript.
- Clarke, C. Manuscript B. Micro structuralism defended via difference making principles. Unpublished Manuscript.
- Fodor, J. (1974). Special sciences (or: The disunity of science as a working hypothesis). *Synthese*, 28, 97–115.
- Fodor, J. (1997). Special sciences: still autonomous after all these years. *Philosophical Perspectives*, 11, 149–163.
- Franklin-Hall, L. R. Forthcoming. High-level explanation and the interventionist's 'variables problem'. *British Journal for the Philosophy of Science*, axu040.
- Garfinkel, A. (1981). *Forms of explanation*. New Haven, CT: Yale University Press.
- Goodman, N. (1954). *Fact, fiction and forecast*. Athlone. Citations refer to 2nd edition (Harvard University Press, 1983).
- Haug, M. C. (2011a). Abstraction and explanatory relevance; or, why do the special sciences exist? *Philosophy of Science*, 78, 1143–1155.
- Haug, M. C. (2011b). Natural properties and the special sciences. *The Monist*, 94, 244–266.
- Hempel, C. G., & Oppenheim, P. (1948). Studies in the logic of explanation. *Philosophy of Science*, 15, 135–175.
- Jackson, F., & Pettit, P. (1992). In defence of explanatory ecumenism. *Economics and Philosophy*, 8, 1–21.
- Jaworski, W. (2002). Multiple-realizability, explanation and the disjunctive move. *Philosophical Studies*, 108, 289–308.
- Kincaid, H. (1986). Reduction, explanation and individualism. *Philosophy of Science*, 53, 492–513. Cited as revised in Kincaid (1997b) chapter 3.
- Kincaid, H. (1987). Supervenience doesn't entail reducibility. *Southern Journal of Philosophy*, 25, 342–356.
- Kincaid, H. (1988). Supervenience and explanation. *Synthese*, 77, 251–281.
- Kincaid, H. (1990). Molecular biology and the unity of science. *Philosophy of Science*, 57, 575–593. Cited as revised in Kincaid (1997b) chapter 4.
- Kincaid, H. (1993). The empirical nature of the individualism-holism dispute. *Synthese*, 97, 229–247. Cited as revised in Kincaid (1997b) chapter 2.
- Kincaid, H. (1996). *Philosophical foundations of the social sciences*. Cambridge: Cambridge University Press.
- Kincaid, H. (1997a). *Defending non-reductive unity. Chapter 5 of individualism and the unity of science*. Lanham: Rowman and Littlefield. Revision of Kincaid (1987) and Kincaid (1988).
- Kincaid, H. (1997b). *Individualism and the unity of the sciences*. Lanham: Rowman and Littlefield.
- Kitcher, P. (1981). Explanatory unification. *Philosophy of Science*, 48, 507–531.
- Kitcher, P. (1984). 1953 and all that. A tale of two sciences. *Philosophical Review*, 93, 335–373.
- Kitcher, P. (1989). Explanatory unification and the causal structure of the world. In P. Kitcher & W. Salmon (Eds.), *Scientific explanation* (pp. 410–505). Minneapolis, MN: University of Minnesota Press.
- Lange, M. (2009). *Laws and lawmakers: Science, metaphysics, and the laws of nature*. Oxford: Oxford University Press.
- Lewis, D. K. (1986). Causal explanation. In *Philosophical Papers* (Vol. 2, pp. 214–40). Oxford: Oxford University Press.
- Little, D. (1993). On the scope and limits of generalizations in the social sciences. *Synthese*, 97, 183–207.
- MacDonald, G. (1985). Modified methodological individualism. *Proceedings of the Aristotelian Society*, 86, 199–211.

- MacDonald, G. (1992). Reduction and evolutionary biology. In D. Charles & K. Lennon (Eds.), *Reduction, explanation and realism* (pp. 69–96). Oxford: Oxford University Press.
- Marchionni, C. (2008). Explanatory pluralism and complementarity: From autonomy to integration. *Philosophy of the Social Sciences*, 38, 314–333.
- Marras, A. (1993). Psychophysical supervenience and nonreductive materialism. *Synthese*, 95, 275–304.
- Owens, D. (1989). Levels of explanation. *Mind*, 98, 59–79.
- Pereboom, D., & Kornblith, H. (1991). The metaphysics of irreducibility. *Philosophical Studies*, 63, 125–145.
- Potochnik, A. (2010). Levels of explanation reconceived. *Philosophy of Science*, 77, 59–72.
- Putnam, H. (1967). Psychological predicates. In W. Capitan & D. Merrill (Eds.), *Art, mind and religion* (pp. 37–48). Pittsburgh, PA: University of Pittsburgh Press. Cited as reprinted in Putnam (1975b).
- Putnam, H. (1973). Reductionism and the nature of psychology. *Cognition*, 2, 131–146. Cited as revised in Putnam (1975c).
- Putnam, H. (1975a). *Mind language and reality: Philosophical papers* (Vol. 2). Cambridge: Cambridge University Press.
- Putnam, H. (1975b). The nature of mental states. In *Putnam 1975a* (pp. 429–441).
- Putnam, H. (1975c) Philosophy and our mental life. In *Putnam 1975a* (pp. 291–303).
- Sawyer, R. K. (2002). Nonreductive individualism: Part I—supervenience and wild disjunction. *Philosophy of the Social Sciences*, 32, 537–559.
- Shapiro, L., & Sober, E. (2012). Against proportionality. *Analysis*, 72, 89–93.
- Sober, E. (1988). Confirmation and law-likeness. *Philosophical Review*, 97, 93–98.
- Sober, E. (1999). The multiple realizability argument against reductionism. *Philosophy of Science*, 66, 542–564.
- Strevens, M. (2008). *Depth: An account of scientific explanation*. Cambridge, MA: Harvard University Press.
- Weatherston, B. (2012). Explanation, idealisation and the goldilocks problem. *Philosophy and Phenomenological Research*, 84, 461–473.
- Weslake, B. (2010). Explanatory depth. *Philosophy of Science*, 77, 273–294.
- Woodward, J. (2003). *Making things happen: A theory of causal explanation*. Oxford: Oxford University Press.
- Woodward, J., & Hitchcock, C. (2003). Explanatory generalizations, Part I: A counterfactual account. *Nous*, 37, 1–24.
- Yablo, S. (1992). Mental causation. *Philosophical Review*, 101, 245–280.
- Ylikoski, P., & Kuorikoski, J. (2010). Dissecting explanatory power. *Philosophical Studies*, 148, 201–219.