

How to define levels of explanation and evaluate their indispensability

Christopher Clarke¹ 

Received: 22 June 2015 / Accepted: 15 February 2016 / Published online: 11 April 2016
© Springer Science+Business Media Dordrecht 2016

Abstract Some explanations in social science, psychology and biology belong to a higher level than other explanations. And higher explanations possess the virtue of abstracting away from the details of lower explanations, many philosophers argue. As a result, these higher explanations are irreplaceable. And this suggests that there are genuine higher laws or patterns involving social, psychological and biological states. I show that this ‘abstractness argument’ is really an argument schema, not a single argument. This is because the argument uses the ‘is lower than’ relation, and this relation admits of different readings. I then suggest four rigorous definitions of the ‘is lower than’ relation, and show that the abstractness argument’s prospects are much brighter for some of these definitions than for others. To show this, I evaluate the so-called ‘disjunctive threat’ to the abstractness argument.

Some explanations belong to a higher level than others. To a very rough approximation, psychological explanations are typically higher than physiological ones, for example; but psychological explanations are typically lower than social ones. And so, for any explanation that correctly explains why a phenomenon occurred, one can ask: is there some lower explanation that also correctly explains the phenomenon? And, if there is, can this lower explanation *replace* the higher explanation? And by ‘replace’ I mean: the higher explanation provides no understanding (of why the phenomenon occurred) over and above the understanding of it that the lower explanation provides.

In practice, philosophers agree that the answer is often that the higher explanation in question can’t be replaced. Suppose—just to explore how absurd the idea is—that every correct higher explanation could indeed be replaced by some correct lower

✉ Christopher Clarke
cjc84@cam.ac.uk

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

explanation. This would in many cases require that the lower explanation describe monstrously complicated lower facts. Instead of describing a nation as being at war, for example, it would need to describe the actions performed by each of several thousand soldiers, civilians and politicians. But human inquirers are not able to entertain such monstrously complicated propositions, let alone communicate such propositions to others. Thus some higher explanations are irreplaceable—at least ‘in practice’ for human inquirers.

Much more controversial, however, is the question of whether some higher explanations are irreplaceable even ‘in principle’.¹ On the face of it, this question seems to be a question about an idealizer inquirer, one who can entertain more complicated propositions. The question asks: do any higher explanations provide this idealized inquirer with understanding above and beyond that provided by every lower explanation? Unfortunately, this question doesn’t specify the extent to which one is supposed, when tackling this question, to idealize away from our cognitive imperfections as humans. Exactly what propositions is an idealized inquirer supposed to be able to entertain? As a result, there is a risk that the controversy over in-principle replaceability will boil down to an insubstantial dispute over the definition of an idealized inquirer.

This paper will clearly identify the substantial issues at stake in this controversy, and will separate them from insubstantial disputes over definitions. Specifically, it will clarify the logic of an argument at the centre of the controversy, which I will call the ‘abstractness argument’; and it will clarify the logic of the so-called ‘disjunctive threat’ to this abstractness argument. For one thing, I will show that the abstractness argument appeals to an ‘is lower than’ relation that admits multiple definitions. This is important because it is easier to neutralize the disjunctive threat to the abstractness argument for some definitions of the ‘is lower than’ relation than for others, I will show. For example, I will show that the abstractness arguments run by several philosophers fail—namely those run by Block, Fodor, Kincaid, Kornblith, Levine, Marras, Pereboom, Putnam, Weslake and at one stage Antony. I hope that the clarificatory work of this paper will allow me to press this criticism more clearly and forcefully than existing criticisms in the literature.

1 The abstractness argument and genuine higher patterns

The abstractness argument normally proceeds by example. One presents a phenomenon for which one has found (i) a relatively high explanation that is (or seems to be) correct, and (ii) one or more relatively low explanations that are (or seem to be) correct. Take, for example, a piece of litmus paper that turned red when dipped into a solution. A relatively high explanation of this fact is (i) that the solution is acidic; and a relatively low explanation of this fact is (ii) that the solution is an HCl solution. (Section 3 will discuss why one might take the latter explanation to be lower than the former.) The argument then continues as follows²: (1) This higher explanation ‘abstracts away’

¹ For example, Kim (1998, pp. 104–105), Rosenberg (2001, p. 756) and Clapp (2001) say that all higher explanations are in principle replaceable. See the citations throughout this paper for those who disagree.

² See footnotes 3–6 and 16–19, excluding Potochnik and Lange.

from some of the details of each of these lower explanations. But by induction: (2) the same holds for every lower explanation that correctly explains why the phenomenon occurred, not just for the lower explanations that were singled out in (ii) above. But (3), when correct, an explanation that abstracts away from some of the details of another explanation provides some understanding that the other explanation does not. It follows that (4) this higher explanation would (if correct) provide some understanding of the phenomenon that no lower explanation provides. (This is not to say that the lower explanations provide no understanding at all of why the phenomenon occurred.)³

The abstractness argument then seems to continue as follows: even if one places no restrictions on the propositions that an (ideal) inquirer can entertain, the above argument for (4) remains sound.⁴ But this suggests that the fact (4)—that some higher explanations are irreplaceable—reflects a feature of the wider world, not just a feature of human cognition. Namely: such higher explanations describe genuine higher laws (Fodor 1974; Pereboom and Kornblith 1991; Kincaid 1996; Baker 2007), or at least genuine patterns amongst higher states.⁵ (And by ‘genuine’ I just mean objectively important or privileged.) In other words, one contends: (5) Genuine higher patterns, if they exist, would constitute an additional reason why some higher explanations are irreplaceable—additional, that is, to reasons arising out of human cognitive limitations. And indeed the considerations in the last paragraph indicate that (6a) the higher explanation in hand describes a genuine higher pattern; whereas (6b) none of the lower explanations describe this higher pattern, at least not as such. (This argument can be formulated in various different ways, depending upon how one understands ‘indicate’ here.)⁶

³ Jackson and Pettit (1992), Kincaid (1997a), Sober (1999) and Marchionni (2008) think lower explanations often can provide understanding of the phenomenon in question; contrast Putnam (1973, pp. 296–298) and Garfinkel (1981). In fact, they think lower explanations can provide some understanding that the higher explanation in question does not; thus the two explanations complement each other.

⁴ Putnam (1973, p. 296), Fodor (1974, p. 113), Kitcher (1984), Kincaid (1997a, p. 89), Sawyer (2002, pp. 540, 552), Woodward (2003, p. 355) and Haug (2011b, p. 245) make it clear that the thesis they defend is not a ‘subjective’ one pertaining to ‘cognitive limitations’ (Haug 2011a). Similarly MacDonald (1985, p. 201), Sober (1999, footnote 9) and Lange (2004, p. 105) defend a thesis about replaceability ‘in principle’. See also Pereboom and Kornblith (1991) and Block (1997). Note for example Jackson and Pettit’s (1992) discussion of the movements of individual molecules in a gas. Clearly this knowledge is only available to highly idealized inquirers; so their defence of their thesis doesn’t rest on any claims about cognitive limitations.

⁵ Putnam (1973, pp. 296–297), Garfinkel (1981, pp. 91–96), Kitcher (1984), Marras (1993, p. 279), Antony (1999, p. 16) as well as Kincaid (1986, pp. 40–43, 1993, p. 24, 1997a) and Potochnik (2010, p. 69) talk of higher explanations ‘capturing’ ‘highlighting’ or ‘bringing out’ patterns.

⁶ *The deductive formulation.* It is quite easy for a pattern to constitute a genuine one: each true explanation describes a genuine pattern. And so true higher explanations describe higher patterns; whereas less abstract lower explanations do not describe these higher patterns as such. In other words, (2) directly entails (6). *The abductive formulation.* It is quite difficult for a pattern to constitute a genuine one. But, whenever a higher explanation (if correct) would provide understanding that the ‘rival candidate explanations’ do not, this provides good evidence that the explanation is true. Indeed, it provides good evidence that the explanation describes a genuine pattern. See for example Sturgeon (1985), Wright (1992, Chap. 5), LaPorte (2004, pp. 19–20), Baker (2007, p. 8), and Beebe and Sabbarton-Leary (2010, Chap. 1). In other words, (4) abductively supports (6). Incidentally, one interesting feature of this ‘inference to the best explanation’ is that the higher explanation and the ‘rival’ lower explanations here are not incompatible explanations; contrast (Lipton 1991).

(It is useful at this point to distinguish the abstractness argument from arguments that appeal to ‘projectability’ (Fodor 1997; Antony 1999). The Projectability Argument: The concepts in higher explanations are often ‘projectible’, in that one is warranted in inductively projecting these concepts from observed cases to unobserved cases. And this indicates that (6a) these higher explanations describe genuine laws. In contrast, the lower explanations that might replace these higher explanations, typically do not use projectible concepts. In this case, (6b) such lower explanations do not describe these higher laws—at least not ‘as such’.⁷ And therefore (4) these higher explanations provide some understanding of the phenomenon in question that no lower explanation provides. This projectability argument constitutes an independent argument for (4) and (6). Or at least it does insofar as projectibility is not closely tied to abstracting away from detail; which it isn’t for Antony (1999) (although contrast Fodor’s projectability argument as examined in Sect. 6). This paper will *not* have anything to say about any independent arguments for (4), such as Antony’s projectability argument.)

At any rate, I will call the thesis that there are genuine higher patterns ‘inflationism’; and I will call its negation ‘deflationism’. And, to interpret the abstractness argument in terms of (1)–(6), as I did above, is to suggest that the controversy over in-principle explanatory dispensability is most fruitfully and clearly understood as, in fact, a metaphysical debate about the truth of inflationism. And this interpretation is a fair one, I contend, because it avoids the risk of the controversy collapsing into a trivial dispute over the definition of an idealized inquirer. This is not to deny, however, that there are also interesting epistemic questions about the various ways in which our human understanding of the world is limited (McGinn 1991). Nor is it to deny that there are alternative ways of understanding in-principle dispensability that will be appropriate for other philosophical contexts (Bedau 2010; Silberstein 2012).

The reason I introduce the terms ‘deflationism’ and ‘inflationism’, rather than sticking with the more familiar ‘reductionism’ and ‘anti reductionism’, is because the latter terms are ambiguous (Silberstein 2012). Often ‘reductionism’ and ‘anti reductionism’ are intended as theses concerning in-principle explanatory replaceability and irreplaceability.⁸ And, on such occasions, I read ‘reductionism’ as deflationism about higher patterns, and ‘anti reductionism’ as inflationism—for the reasons I’ve just given. But on other occasions ‘reductionism’ means the thesis (a) that there are ‘bridge laws’ between higher theories and lower theories; or (b) that lower laws or facts not only ‘fix’ the higher laws and facts, but also ‘ground’ them metaphysically or explanatorily (Correia 2012); or even (c) that each higher property or causal power or kind is identical to some lower property or causal power or kind. The problem is that some philosophers are reductionist in senses *a b* or *c*, but are nevertheless inflationists about higher patterns. For example, Antony (2003, p. 14), Sober (1984; 1999) and perhaps McGinn (1991) think that some higher patterns are genuine; but that nevertheless all

⁷ Antony (2008) contends that higher properties are identical to lower properties. But lower explanations fail to describe higher laws *as such*, because they use lower concepts.

⁸ Indeed this is the original meaning of the terms. See Oppenheim and Putnam (1958), Garfinkel (1981, pp. 15, 44), MacDonald (1985, p. 202), Marras (1993), Kincaid (1997b, e.g. 3), Antony (1999, p. 18) and Sober (1999).

higher patterns are in a sense grounded in lower patterns. In other words, they think that the genuineness of a higher pattern consists in something other than it being a brute pattern, ungrounded in lower patterns. Indeed, Antony and Sober think that a description of a higher pattern can be explanatorily irreplaceable in principle, despite the fact that the higher pattern itself can be explained by lower facts. Therefore, to allow room for such hybrid positions, I prefer ‘deflationism’ and ‘inflationism’ to the ambiguous terms ‘reductionism’ and ‘anti reductionism’.

2 The disjunctive threat to the abstractness argument

There are many concerns one might have regarding the abstractness argument. (I am suspicious, for example, of the move from premise four to premise six.) This paper will limit itself, however, to examining the inductive step between premise one and premise two. Both premise one and premise two say that a given higher explanation abstracts away from some lower explanations. But premise one concerns only those lower explanations that have been explicitly singled out for illustration; whereas premise two concerns each and every lower explanation of the phenomenon in question. The latter include lower explanations that are very complicated, or that are otherwise obscured or hidden. This section will describe the ‘disjunctive threat’ to this inductive step. By analyzing the disjunctive threat more thoroughly than existing treatments, this section will show that there are exactly five ways in which one might try to neutralize it.

Here goes. Suppose NECESS: the lower states of an entity ‘fix’ its higher states. To be precise, suppose that it’s possible for some entity to be simultaneously in higher state H (being an acidic solution) and in lower state L (being an HCl solution). Then NECESS say it’s absolutely necessary that, if any entity is in lower state L , then it is also in higher state H . A few comments. (1) Throughout this paper, when I talk of ‘states’ I will mean general state types (being an acidic solution) rather than particular state tokens (by next week this wine will be an acidic solution). (2) By saying that p is ‘absolutely necessary’ I mean that if p were not the case, it would be absurd. More formally: ‘the necessary is that whose negation counterfactually implies a contradiction’ (Williamson 2007, p. 157).

But it follows from the above, everyone recognizes, that each higher state is ‘Kim-parallel’ to some lower state, or to some ‘compound’ thereof. To illustrate this, take again the state of being an HCl solution, and grant that it’s absolutely necessary that any entity in this lower state will be an acidic solution; if only for the sake of illustration. Now consider a second lower state type for which this is also true, for example: the state of being an H_2SO_4 solution. Indeed consider a third lower state type—and so on perhaps to infinity—until one has explicitly numbered all the (possible) lower state types for which it’s absolutely necessary that an entity in that state is an acidic solution. Now define state S as the state of being either in the first lower state (HCl) or in the second (H_2SO_4) or in the third, and so on. So, necessarily, an entity is in state S if and only if it is one of these lower states. (I will say that this definition uses the above disjunction to define S as a ‘compound’ of these other states.) It follows that

it's absolutely necessary that any given entity is an acidic solution if and only if it's in state *S*.⁹

This is what I mean by saying that the state of being an acidic solution is Kim-parallel (K-parallel) to *S*: it is absolutely impossible for something to be in the one state, without it also being in the other; and vice versa. (I should emphasise that there is no assumption here that state *S* constitutes a genuine property.)

Now suppose INFINITE: an idealized inquirer can grasp infinitely long disjunctions of concepts. Therefore she can grasp a concept *S* 'being either in the first lower state (HCl) or the second (H₂SO₄) or the third, and so on'; where these states are explicitly named as above. In other words, concept *S* is K-parallel to the concept of being an acidic solution. As a consequence, an idealized inquirer can take any explanation that uses relatively high concepts (being an acidic solution) and replace each of these concepts with the K-parallel concept (*S*). This generates what I will call a K-parallel explanation.

Finally, suppose SPLIT: concepts that are grasped in an explicitly disjunctive form are really composed of several concepts in combination. Concept *S*, for example, is really composed of a concept that refers to lower state one (HCl) and a concept that refers to lower state two (H₂SO₄) and so on; as above. And suppose also STATES: any explanation that refers only to lower states counts as a lower explanation. It follows that K-parallel explanations are lower than the explanations out of which they are constructed.

(However, what if we suppose instead that SPLIT is false? That is to say, concept *S* is instead a single concept that refers to compound state *S*. Let's suppose COMPOUND, however: any compound of lower states is itself a lower state. So *S* is a lower state. It follows again, given STATES, that K-parallel explanations are indeed lower.)

Let's take stock: if (a) INFINITE holds, and if (b) NECESS holds, and if (c) either SPLIT or COMPOUND holds, and if (d) STATES holds, then for each higher explanation there will be a lower K-parallel explanation.

This raises an important worry: why think that what holds for typical lower explanations will also hold for these lower K-parallel explanations? Therefore, why trust the inductive inference in the abstractness argument: the singled-out lower explanations are less abstract than the higher explanation in question, therefore all lower explanations—including lower K-parallel explanations—are less abstract than it. This is the 'disjunctive threat' to the abstractness argument.

This section has shown that there are exactly five ways to neutralize the threat: (A) deny that INFINITE holds; (B) deny that NECESS holds; (C) both deny that SPLIT holds and deny that COMPOUND holds; (D) deny that STATES holds; or (E) show explicitly that K-parallel explanations are less abstract than the higher explanations out of which they are constructed. (It is important to note that the disjunctive threat, as I've characterized it, has a modest aim: it seeks to undermine the abstractness argument for inflationism. It doesn't itself constitute a positive argument for deflationism. Nor does it pose a threat to Antony's projectability argument, for example. So, although one might attempt to recast the disjunctive threat

⁹ See Kim (1984, pp. 169–70), Bacon (1986) and Kim (1991, p. 151).

in more wide-reaching terms, this may expose it to further objections, in addition to *A–E*.)

Response *A* (the rejection of INFINITE) can be quickly dismissed. This response insists that there are restrictions on what an idealized inquirer can do. But any inflationist who insists on such restrictions would thereby trivialize the debate over ‘in principle’ explanatory dispensability—making it an insubstantial dispute over the definition of an idealized inquirer (Sect. 1). More importantly, such a restriction would make it difficult for inflationists to draw any substantive metaphysical conclusions—in particular inflationism about higher patterns—from the fact that a higher explanation cannot be replaced by any lower explanation (Sect. 1). Perhaps for these reasons, few inflationists deny INFINITE, even in light of the disjunctive threat; and many explicitly accept it.¹⁰

Response *B* (reject NECESS) is more promising. This response says that some higher states aren’t fixed by lower states. One might, for example, endorse NECESS only as a claim about nomic necessity—the necessity associated with the laws of nature—not absolute necessity. But this would be tantamount to rejecting the mainstream inflationist position, in favour of a more ‘emergentist’ inflationism.¹¹ Indeed, of the inflationists who endorse the abstractness argument, very few explicitly deny NECESS as a claim about absolute necessity, even in light of the disjunctive threat.¹² So any inflationist who wants to endorse response *B* should, at least, be more upfront about it. For this reason, I won’t examine this response further.

With response *A* to the disjunctive threat dismissed, and response *B* set aside, the rest of the paper will examine the prospects for responses *C*, *D* and *E*. Section 3 identifies four rough ways to define the distinction between higher and lower states, and sharpens them up. On this basis, it shows how the viability of response *C* depends upon which definition of the ‘is lower than’ relation one has in mind here. Section 4 draws four corresponding distinctions between higher *explanations* and lower explanations. And it examines response *D* in light of this. Section 5 clarifies the abstractness argument further by examining various conceptions of abstractness. On this basis, Section 6 argues that response *E* fails. Section 7 concludes.

3 Higher and lower states

What do philosophers mean when they say that a given state is ‘lower than’ another? Take, for example, the state *A* of being an acidic solution. Some take this state to be essentially dispositional: necessarily, acidic solutions tend to cause alkalis to react, litmus paper to turn red, and so on. (More rigorously: necessarily, solutions in state *A* are in some state *X* that tends to cause alkalis to react, litmus paper to turn red, and so

¹⁰ See Kincaid (1986, p. 38, 1993, p. 23), Weslake (2010, pp. 287–289) and especially Kincaid (1997a, pp. 72–74) for an explicit endorsement. It’s also very clear that Fodor (1974, 1997), Pereboom and Kornblith (1991), Block (1997), and Antony (1999) endorse INFINITE too. Refer also to Sect. 6 and to footnote 4.

¹¹ See van Cleve (1990, p. 222), Chalmers (2006) and Noordhof (2010, pp. 71–73) for this characterization of the mainstream position and of emergentism; and Silberstein (2012) for complications.

¹² See all the inflationists cited in Sects. 1, 6 and the present section.

on.)¹³ Thus the state of being acidic is necessarily associated with the above ‘causal role’. Contrast this state with those states that ‘occupy’ this causal role, for example the state of being an HCl solution.¹⁴ In virtue of this, some would count the state of being an HCl solution as lower than the state of being an acidic solution.¹⁵ In other words, the ‘is lower than’ relation is sometimes explicitly defined as the ‘occupies the role of’ relation; especially by philosophers of mind.

But other philosophers do not give the ‘is lower than’ relation a rigorous definition. Indeed, the literatures in general philosophy of science, and the philosophy of biology, economics and the social sciences gesture towards three alternative ways of defining this relation in addition to the definition above, I suggest. And some philosophers fail to appreciate that the other philosophers whom they cite, and with whom they engage, understand this relation differently from how they themselves understand this relation.¹⁶

Alternative One: Specificity. One might imagine that it’s absolutely impossible (in the sense given in Sect. 2) for an HCl solution to fail to be an acidic solution; but not vice versa. Thus the former state is, in this sense, more ‘specific’ than the latter. Or one might think, at least, that the laws of nature entail that if an entity is an HCl solution, then it is also an acidic solution; but not vice versa. Thus the former state is more specific than the latter, in a ‘nomic’ sense. In virtue of this, some philosophers seem to count the former state as lower than the latter. In other words, the ‘is lower than’ relation is taken to be either the ‘is absolutely more specific than’ relation or the ‘is nomically more specific than’ relation. Although this definition is usually endorsed only implicitly (Teller 1984; Sober 1999; Strevens 2008), on rare occasions it is made more explicit¹⁷:

In the case of simple event explanations at different levels, the values of the high-level variable will be coarser than those of the low-level variable. . . . Such is the relationship between cause variables cited in the following two explanations of a crash: First, a low-level explanation appealed to a binary variable, one of whose values was the exact speed of the car as it turned a bend, 50 mph, and another value representing some speed at which the car would not have crashed, such as 20 mph. Second, a high-level explanation appealed to another binary variable, with one value corresponding to the car’s speed exceeding 30 mph and the other value corresponding to the car not exceeding this speed. (Franklin-Hall Forthcoming, Sect. 3)

¹³ Note that this leaves open the question of whether $X = A$; in constant to $X =$ ‘being an HCl solution’, for example.

¹⁴ See Lewis (1972) and Shoemaker (2007, Chap. 2) for the ‘lower order’ and the ‘subset’ accounts of role occupation respectively.

¹⁵ See Lewis (1972), Fodor (1974), Shapiro (2000), Shoemaker (2007), and Haug (2011a).

¹⁶ See the footnotes throughout this section. For example, Weslake (2010, p. 274) and Lange (2004, p. 105) each assume that Jackson and Pettit (1992) adopt the same understanding of the relation as they do. But this is probably false; see Sect. 4 for Lange, and footnotes in the present section for Weslake, Jackson and Pettit.

¹⁷ This is also Jackson and Pettit’s (1992, pp. 6–7) official definition; although they don’t always stick to it (see following footnotes).

Some comments. (1) The ‘is more specific than’ relation and the ‘occupies the role of’ relation will not issue in anything close to a total ordering of states from lower to higher.¹⁸ Thus these relations do not partition the class of states into discrete levels. (2) The latter relation only applies when the higher state is essentially dispositional. But the former relation also applies to states that are associated with a particular causal role only contingently. So the former relation may come apart from the latter; see also Haug (2011a).

Alternative Two: The Microscopic. Consider the state of being ionized. This state can be instantiated by relatively small entities, such as an H^+ particle. And it is an intrinsic state of such entities. Contrast this with the state of being an acidic solution, which can be instantiated by larger entities *only*—by a body of liquid, for example. Also contrast the state of being ionized with the state of being the fastest moving ion in a solution, which is not an *intrinsic* state of any relatively small entity. On this basis, let’s call the state of being ionized a more ‘microscopic’ state than these other two states. In virtue of this, some philosophers seem to count the state of being ionized as lower than both these states. Thus the ‘is lower than’ relation is taken to be something like my ‘is more microscopic than’ relation. Although this definition is usually endorsed implicitly, on rare occasions it is made more explicit¹⁹:

The explanations in Garfinkel’s rabbit example are formulated at lower and higher levels in the sense that one deals with an individual and the other with a population of individuals. In general, a lower-level explanation cites properties of objects that stand in a part–whole relationship to objects referenced in the competing higher-level explanation. (Potochnik 2010, p. 64)

Some comments: (1) I’ve used the ‘is smaller than’ relation between entities as a basis for defining the ‘is more microscopic than’ relation between states. An alternative is to instead use the ‘is a part of’ relation between entities.²⁰ (2) Another alternative is to replace the intrinsicity restriction with the restriction that the state constitute a genuine property. (3) This definition applies only to unary states. But what about binary states, such as the state of x exerting an intermolecular van der Waals force on y ? I count this state as relatively microscopic. This is in virtue of the fact that it can be instantiated by a pair of relatively small entities, namely molecules; and the fact that is an intrinsic state of the pair.²¹

¹⁸ On these two definitions we have: $[HCl] < [Acid]$; and $[Acid] \leq [150J \text{ Mean kinetic energy}]$; and $[150J \text{ Mean Kinetic Energy}] < [Hot]$. (Read $<$ as ‘is lower than’ and \leq as ‘is not higher than’.) Suppose for reductio that $<$ constituted a total ordering. It would follow that $[HCl] < [Hot]$. But $[HCl]$ is not lower than $[Hot]$, on either of these two definitions. So we have a contradiction. So $<$ is not a total ordering.

¹⁹ Zahle (2003) and Marchionni (2008, p. 316) explicitly endorse something roughly like this; although they don’t discuss details involving intrinsic states and binary relations. And Weslake (2010, p. 287) is implicitly committed to something like this, even though it sometimes looks as if he has the ‘is more specific than’ relation in mind instead. Garfinkel (1981) and Ylikoski (2011) may perhaps have a hybrid of the ‘is more specific than’ and ‘is more microscopically structured than’ relation in mind. It’s unclear. At points Jackson and Pettit (1992) seem to adopt such a hybrid, although this is not their official view (see footnotes above).

²⁰ See Potochnik and McGill (2012) for some potential shortcomings of this approach.

²¹ See Weatherson and Marshall (2014, Sect. 1.3) on intrinsic binary relations.

Alternative Three: Microscopic Structure. The state of being an HCl solution is the state of containing a high number of H^+ and Cl^- ions amongst a high number of H_2O molecules, where the intermolecular bonds are of moderate strength. Since this state cannot be instantiated by any relatively small entity (a molecule) it doesn't count as a relatively microscopic state. It is, however, a state of 'being such that' some relatively small entities (molecular parts of the solution) or pairs thereof instantiate a particular 'permutation' of microscopic states, unary or binary. Contrast this with the state of being an acidic solution, which some philosophers suppose is not such a state.²² I will assume for the moment that this supposition is correct, just for the sake of illustration. On this basis, let's say that the state of being an HCl solution is 'more microscopically structured' than the state of being an acidic solution. In virtue of this, some seem to count the state of being an HCl solution as lower than the state of being an acidic solution. Thus the literature gestures towards defining the 'is lower than' relation as something like my 'is more microscopically structured' relation. Kim (1998, p. 82) is more explicit on this point when he talks of:

micro-based (or microstructural) properties, properties of a whole that are characterized in terms of its microstructure.

A few clarificatory comments. (1) The state of 'being either an HCl solution or an NaCl solution' is a microscopically structured state. But note that this state is not specific enough to fix whether or not an entity is acidic. Thus, when defining the 'is lower than' relation as the 'is more microscopically structured' relation, one needs to modify NECESS. Namely: for each higher state, there is *some* 'partition' of lower states, each of which fixes whether or not the higher state in question obtains.²³ (2) By 'permutation' I mean a Boolean permutation, such as conjunction or disjunction for example (Baker 1993, p. 81). (3) I have defined permutations as Boolean permutations of unary and binary states; but not of tertiary or quaternary states or the like. Otherwise some philosophers may worry that for their purposes it's too easy for states to count as microscopically structured. (4) I do not place any spatio-temporal restrictions on permutations (Noordhof 2010, p. 72). So the state of 'being the most concentrated HCl solution ever' counts for example as a microscopically structured state, albeit a highly extrinsic one. It's the state of being such that no other solution anywhere—past, present or future—contains a greater ratio of H^+ and Cl^- ions to H_2O molecules. (5) Accordingly, NECESS becomes equivalent to: no difference in a higher state of a given entity, without a difference in the microscopic states, unary or binary, of *some* entity *somewhere* in the world, *at some point* in the past, present or future (see Sect. 2).

(6) One might think, contrary to what I supposed above, that the state of being acidic is as microscopically structured as the state of being an HCl solution (Mumford 1998). So the 'is more microscopically structured' relation may well come apart from the 'occupies the role of' relation. (7) 'Being an HCl solution' and 'being an acidic solution' are states instantiated by entities of the same size (bodies of liquid). So

²² For example, Prior et al. (1982) claim that essentially dispositional states—such as the state of being acidic—are distinct from the microscopically structured properties in which they are grounded.

²³ By 'partition' I mean that a collection of states for which: necessarily, the entity in question will be in one of the states in the collection.

neither of these two states is more microscopic than the other. This fact illustrates how the occupant–role, specific–general, and ‘micro structured’–‘macro structured’ relations come apart from the microscopic–macroscopic relation. (8) This same fact also shows that the argument in Sect. 2 cannot be run directly when ‘is lower than’ is defined as ‘is more microscopic than’.²⁴

In sum, the literature gestures towards four ‘is lower than’ relations, and these four characterizations can be tightened up as above.

This puts us in a position to evaluate COMPOUND, which says that any compound of lower states counts as a lower state itself. Firstly, if several states are each an intrinsic state of relatively small entities, then the compound of these states will be as well. In other words, any compound of microscopic states is itself microscopic. So, on the microscopic definition of lower states, COMPOUND is true. Secondly, any compound of Boolean permutations of microscopic states is itself a Boolean permutation of microscopic states. In other words, any compound of several microscopically structured states is itself a microscopically structured state. So, on the microscopically structured definition of lower states, COMPOUND is true. Thirdly, some compounds of specific states are not themselves relatively specific. For example, the state of being either light blue or dark blue is less specific than the state of being light blue, and also less specific than the state of being dark blue. Thus, on the definition of lower states as relatively specific states, COMPOUND is false (Teller 1984). Fourthly, a much longer discussion would be needed on the question of whether COMPOUND holds for the definition of lower states as occupants of roles. The answer will depend upon controversial issues concerning causation (List and Menzies 2009) and property identity. To sum up: is response *C* to the disjunctive threat—which requires rejecting COMPOUND—a viable response? The answer depends upon how one defines the ‘is lower than’ relation.

4 Levels of explanation

Each of these four ‘is a lower state than’ relations generates a corresponding distinction between relatively high explanations of a phenomenon and relatively low explanations thereof. To illustrate this, let’s ask: why did Alexander Litvinenko die? For this explanandum, compare each explanans on the following list. Explanans *A*: Many of Litvinenko’s cells contained Polonium 210. Explanans *B*: Andrey Lugovoy wanted Litvinenko dead. Explanans *C*: Putin adopted an institutional policy of murdering dissidents.

I will assume for illustration that explanans *A* refers to exactly one state, a state that is necessarily instantiated if and only if a given entity contains Polonium 210. Similarly, explanans *B* refers to exactly one state, a state that is necessarily instantiated if and only if a given entity wants Litvineko dead. (Of course, *A* and *B* don’t refer to states only; they also refer to individual entities, such as Litvineko and Litvineko’s cells.) But relatively small entities (cells) can contain Polonium 210, whereas only larger

²⁴ Instead one has to proceed indirectly: Sect. 2 provided an argument that for each macroscopic explanation there is a K-parallel explanation that only refers to *microscopically structured* states. It follows immediately that for each macroscopic explanation there is a K-parallel explanation that only refers to *microscopic* states.

entities (multi-cellular organisms) can want Litvinenko dead. So the state to which *A* refers is more microscopic than the state to which *B* refers. But, for the sake of illustration, let's take the 'is lower than' relation to be the 'is more microscopic than' relation. Therefore explanans *A* offers a lower explanation than *B*, one might say.

What about *B* and *C*? The state of wanting Litvinenko dead is an intrinsic state of a multi-cellular organism. Or so one might suppose for illustration.²⁵ In contrast, adopting an institutional policy of murdering dissidents is an intrinsic state only of a larger entity, such as the Kremlin. It is not an intrinsic state of anything organism-sized or smaller. (After all, for Putin to adopt an institutional policy, there needs to be a large institutional apparatus to begin with.) Thus the state to which *B* refers is lower than the state to which *C* refers. Therefore explanans *B* offers a lower explanation than *C*, one might say.

In sum, although levels of explanation are rarely rigorously defined, the foregoing discussion illustrates one way of doing so in terms of states. Namely: lower explanations refer to lower states. More carefully:

(STATES—Rigorous) Consider explanans *L* and explanans *H*. Let *l* denote the states that *L* refers to, and *h* denote the states that *H* refers to. And let *l-h* denote the states that *L* refers to, excluding those that *H* refers to. Define *h-l* similarly.

(1) Suppose that there are some states in *l-h* that are lower than some states in *h*; but that no states in *l-h* are higher than any states in *h*. Then define *L* as a lower explanation than *H*. (2) Suppose that there are some states in *h-l* that are higher than some states in *l*; but that no states in *h-l* are lower than any states in *l*. Then define *L* as a lower explanation than *H*. (3) Otherwise, define *L* as not lower than *H*.

When defining lower states as microscopic states, one may wish to add the following restriction: lower explanations attribute states to small entities only (cells). They may refer to large entities (Litvinenko), but only in order to refer to their parts (Litvinenko's cells).

STATES offers a well-motivated definition of levels of explanation, no matter which of the four ways of distinguishing higher states from lower states one has in mind. Firstly, contrast the explanation 'the solution was acidic and it had a high temperature' with 'the solution was an HCl solution and it had mean kinetic energy of 150J'. The latter is paradigmatically a lower explanation. Nevertheless, being an HCl solution is *not* more specific than, and does not occupy the role of, having a high temperature. So some of the states to which this higher explanation refers are not higher than every one of the states to which this lower explanation refers, one might think.²⁶ And STATES respects this. Secondly, (1) ensures that adding a relatively low state to an explanation makes the explanation lower. And (2) ensures that adding a relatively high state to an explanation makes the explanation higher. Neither clause is redundant.

²⁵ Read this mental state in a suitably 'internalist' way. The story for 'externalist' mental states will be very different.

²⁶ Contrast Marchionni (2008) who has a more restrictive definition in mind, I suspect.

Although STATES is implicitly presupposed by almost all inflationists in the debate, there is a single but noteworthy exception.²⁷ Lange (2004) takes lower explanations to be those that invoke lower counterfactuals rather than higher counterfactuals. But he doesn't define lower counterfactuals as those that relate lower states to each other. Instead Lange has his own (highly intricate) account of whether a counterfactual belongs to a lower science like fundamental physics, as opposed to a higher science like sociology. Thus Lange is implicitly committed to denying STATES. But this point goes unnoticed in the literature. Even Lange himself doesn't mention it. And so the existence of response *D* to the disjunctive threat goes unrecognised, namely: deny STATES in such a way that K-parallel explanations do not count as lower than the explanations out of which they are constructed.

At any rate, I take the question of whether STATES is true to be a simple matter of stipulation. So it suffices to leave the discussion here.

5 Higher explanations as abstract

The last section examined the difference between higher explanations and lower explanations. The task now is to consider in what respects part of an explanation might 'abstract away from some of the details' found in part of another explanation. To this end, I will draw on the work of Clarke (Forthcoming), which outlines four such respects. (The reasons why some philosophers have taken each of these varieties of abstractness to be explanatorily virtuous are complex; see Clarke (Forthcoming) for a detailed discussion.)

Suppose, for illustration, that a given piece of litmus paper turned red when it was dipped into a solution. What explains this? Explanation *H*: 'The solution was an acidic solution. Indeed, consider all the possible types of acid, and all the possible (non-negligible) concentrations of these acids. If any solution were of any of these types and any of these concentrations, the litmus paper would turn red when dipped into it.' Explanation *L*: 'The solution was an HCl solution. (Minor variant: the solution contains a large number of H⁺ and Cl⁻ ions amongst a large number of H₂O molecules.) Indeed, consider all the possible (non-negligible) concentrations of HCl. If any HCl solution were of any of these concentrations, the litmus paper would turn red when dipped into it.'

In what senses, if any, might *H* be more or less abstract than *L*?

Firstly, *H* appeals to generalization *h*: 'for all solutions *x*, and for all permutations λ of acid type and concentration, if the solution *x* were in state λ , the litmus paper would turn red.' But consider the class of individuals over which the individual variable *x* ranges in generalization *h*, namely the class of all solutions. Clarke calls this the 'primary jurisdiction' of generalization *h*. Compare this to the generalization to which *L* appeals, *l*: 'for all solutions *x*, and for all concentrations of HCl λ , if the solution

²⁷ Several deflationists, it should be noted, would deny STATES. These deflationists say that all explanations refer to lower states. What distinguishes higher explanations is that higher explanations from lower ones refer to lower states by using higher *concepts*, they say. Note that on this view the lower state that a given higher concept refers to will vary from semantic context to semantic context; see Kim (1998), Heil (2003) and in some moods Lewis (1994).

x were in state λ , the litmus paper would turn red.’ Here the individual variable x also ranges over the class of all solutions. And, in virtue of their having the same primary jurisdiction, some philosophers would count the two generalizations as equally abstract, Clarke notes.

Secondly, note the class of states over which the *predicate* variable λ ranges in generalization l , namely all concentrations of HCl. Clarke calls this the ‘secondary jurisdiction’ of generalization l . Contrast this with the class of states over which λ ranges in generalization h , namely all permutations of acid type and concentration. The former is a subclass of the latter. Thus h has a wider secondary jurisdiction. In virtue of this, some philosophers would count generalization h as more abstract than l , Clarke notes. See Clarke (Forthcoming, Sect. 6) on the relationship between a generalization’s primary and secondary jurisdictions.

Thirdly, consider the particular factors that an explanation invokes, as opposed to the generalizations it invokes. H for example says that the solution ‘is an acidic solution’, whereas L says that the solution ‘is an HCl solution’. Some philosophers, Sect. 3 observed, consider the latter factor to be more specific than the former, in a nomic or even an absolute sense. The former is more nomically or absolutely general. In virtue of this, some philosophers would count the former factor as more abstract than the latter, Clarke notes.

Fourthly, notice that H is somewhat more syntactically simple than the minor variant of L , which talks of a solution ‘containing a large number of H^+ and Cl^- ions amongst a large number of H_2O molecules, with moderate intermolecular forces’ whereas H talks only of a solution being ‘acidic’. Because of this, H is somewhat easier to grasp than this variant of L : it is more ‘cognitively transparent’, as it were. And in virtue of its relative cognitive transparency, some philosophers would count H as abstracting somewhat away from the details of this variant of L , Clarke notes.

In sum, I’ve sketched four respects in which philosophers have thought that part of an explanation can abstract away from some of the details found in part of another explanation. And, in the case of H and L , it turns out, there are three respects in which the higher explanation H abstracts away from lower explanation L . Of course, it’s possible that the same will not be true for all pairs of higher and lower explanations: some pairs of higher and lower explanations may not be analogous to H and L ; a point to which I now turn.²⁸

6 The abstractness of lower K-parallel explanations

Recall response E to the disjunctive threat. This response concedes that for each higher explanation there will be a lower K-parallel explanation, but it argues that (K-ABSTRACT) the higher explanation in question abstracts away from some of the details of its lower K-parallel explanation. Therefore the inductive step in the abstractness argument survives the disjunctive threat (Sect. 2), it is claimed.

This section will argue that response E fails. Now, some readers will agree with me here, but will think that E ’s failure is so obvious that to argue against E would be to

²⁸ See also Block (1995) and Potochnik (2010) for some very different examples.

argue pointlessly against a straw man. Not so. Many advocates of the abstractness argument very clearly concede that for each higher explanation there is a lower K-parallel explanation: Block, Fodor, Haug, Kincaid, Kornblith, Levine, Marras, Pereboom, Putnam, Weslake (and at one stage Antony).²⁹ These inflationists take themselves to have a substantive disagreement with deflationists who advocate ‘reduction via disjunction’ (Kincaid 1997a, p. 73), where this ‘disjunction strategy’ (Antony and Levine 1997, p. 87) involves replacing higher concepts with ‘open-ended disjunctions’ (Pereboom and Kornblith 1991) or ‘heterogeneous physical–chemical disjunctions’ (Block 1997). In other words, this disagreement is presented as a substantive disagreement about whether lower K-parallel explanations can replace the higher explanations out of which they are constructed—rather than an insubstantial dispute over whether K-parallel explanations count as lower, or whether K-parallel explanations are available to idealized agents.

This section will argue that K-ABSTRACT is false. No explanation abstracts away from its lower K-parallel explanation. I will demonstrate this for each of the standard conceptions of abstractness in turn, namely those delineated in Sect. 5. This is an important task, because many advocates of the abstractness argument do not indicate clearly which of these conceptions of abstractness they have in mind.³⁰

Breadth of a Generalization’s Jurisdiction. Fodor (1997, pp. 155–158) focuses on the generalizations that an explanation invokes. And he understands abstractness in terms of the breadth of jurisdiction of these generalizations (recall Sect. 5). He runs the abstractness argument accordingly, and bolsters it against the disjunctive threat by maintaining K-ABSTRACT: some explanations abstract away from their lower K-parallel explanations.³¹ Working with the present interpretation of abstractness, K-ABSTRACT becomes: some generalizations have a broader primary (or secondary) jurisdiction than their lower K-parallel generalizations. However, every lower K-parallel generalization is generated by replacing the concepts in a generalization with the lower K-parallel concepts. So K-ABSTRACT requires that some concepts apply to more possible individuals (or states) than their lower K-parallel concepts. But this requirement is never met: any concept and its lower K-parallel concept will be necessarily co-extensional; see Sect. 2. Therefore K-ABSTRACT is false, on the present interpretation of abstractness as breadth of primary or secondary jurisdiction.

Contrast my objection to Fodor’s reponse with Weslake’s (2010) objection to it. Weslake contends that typical lower explanations invoke generalizations that have

²⁹ See the extensive citations throughout Sects. 1, 2 and the present section; in particular Fodor (1974; 1997), Kincaid (1986, p. 38, 1997b, p. 3), Antony and Levine (1997), Antony (1999), Weslake (2010, pp. 287–288), Haug (2011a, 2011b) and especially Kincaid (1997a) and perhaps also Putnam (1967, p. 437). Consult Pereboom and Kornblith (1991) and Block (1997), who endorse response *E* themselves, and also attribute it to Putnam and Fodor.

³⁰ Their vague idea is that higher explanations are ‘essentially involved’ in ‘capturing’ higher patterns which appear ‘arbitrary’ and ‘heterogeneous’ from the lower ‘point of view’. This is how Kitcher (1984), Kincaid (1986, p. 34), Pereboom and Kornblith (1991), and Antony and Levine (1997, p. 94) express the idea.

³¹ Fodor argues that breadth of jurisdiction indicates ‘projectability’, which in turn indicates lawhood. But it is clear that lawhood is tied very closely to explanation for Fodor (1997, p. 149). So it is fair to interpret Fodor as running a version of the explanatory abstractness argument; as I do, and as Sober (1999, footnote 17) does.

a broader jurisdiction than those invoked in typical higher explanations. Think for example of Newton's laws in contrast with the principle of natural selection. But this contention is not uncontroversial (Cartwright 1983; Block 1995). One nice feature of my objection to Fodor's response is that I remain neutral on this issue; thus my objection is thereby dialectically more powerful than Weslake's. Another dialectical strength of my argument is that it didn't have to deny that generalizations with broader jurisdictions provide additional understanding; in contrast to Waters (1990, pp. 132–36) and probably Elster (1989).

General Factors. Block (1995, Sect. 3.3), Weslake (2010) and perhaps Putnam (1967, p. 437) and Kincaid (1990, p. 63) seem to focus instead on the particular factors that an explanation invokes, not the generalizations that it invokes.³² And they seem to understand abstractness in terms of the absolute or nomic generality of these factors (recall Sect. 5). They run the abstractness argument accordingly, and bolster it against the disjunctive threat by maintaining K-ABSTRACT: some explanations abstract away from their lower K-parallel explanations. On the present interpretation of abstractness, K-ABSTRACT becomes: some explanations invoke particular factors that are more absolutely / nomically general than the factors that their lower K-parallel explanations invoke. However, every lower K-parallel explanation is generated by replacing the concepts in an explanation with the lower K-parallel concepts. So K-ABSTRACT requires that some concepts are more absolutely / nomically general than their lower K-parallel concepts. And this in turn requires that something's falling under a concept does not entail that it fall under the lower K-parallel concept. But this requirement is never met: any concept and its lower K-parallel concept will be necessarily co-extensional; see again Sect. 2. Therefore K-ABSTRACT is false, on the present interpretation of abstractness as invoking general factors rather than specific ones.

Cognitive Transparency. Marras (1993, p. 284) and Kincaid (1997a, p. 83) seem to endorse the abstractness argument with abstractness understood as cognitive transparency. And they bolster it against the disjunctive threat by maintaining K-ABSTRACT: some explanations abstract away from their lower K-parallel explanations. Working with the present interpretation of abstractness as cognitive transparency, K-ABSTRACT becomes: some explanations make the knowledge that is relevant to the explanation more cognitively transparent than their lower K-parallel explanations do.

It is important to realize however that cognitive transparency is not an absolute. Instead it's relativized to an inquirer's cognitive abilities. To take an example from Clarke (Forthcoming): to a layperson the mathematical expression $Q > W$ counts as much more transparent than $(\sqrt{Q} + \sqrt{W})(\sqrt{Q} - \sqrt{W}) > 0$. But to a highly-trained mathematician, it is immediately apparent that these expressions are mathematically equivalent. Next, recall from Sect. 1 that what is probative in the discussion between deflationists and inflationists is the understanding possessed by an idealized inquirer. But to an idealized inquirer all explanations will be equally cognitively transparent. Therefore both explanations here count as equally abstract for present purposes. So

³² Haug (2011a) may also interpret Fodor (1997) this way, but I'm skeptical of this interpretation.

K-ABSTRACT is false, on the present interpretation of abstractness as cognitive transparency.

At this point, one might object to my reading of Marras and Kincaid: what Marras and Kincaid are really claiming is that some explanations are more *syntactically simple* than their lower K-parallel explanations; and this syntactic simplicity provides additional understanding in itself, they contend, independent of its benefits for cognitive transparency. But this latter contention is false, I insist. Explanatory 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 consider two idealized inquirers who use the $Q > W$ expression and the $(\sqrt{Q} + \sqrt{W})(\sqrt{Q} - \sqrt{W}) > 0$ expression in their respective explanations of some phenomenon. In this case, the two explanations say the same things about the world: they are logically equivalent and therefore synonymous. And the two explanations are cognitively processed with equal ease, because the inquirers are idealized. It follows that neither explanation provides understanding that the other does not provide. This shows that syntactic simplicity provides no additional understanding in itself, independent of its benefits for cognitive transparency.³³ So the objection fails.

In summary, many inflationists concede that for any given higher explanation there is a lower K-parallel explanation. To neutralize the threat that this concession poses to the abstractness argument, these inflationists promise to establish K-ABSTRACT: some explanations abstract away from some of the details of their lower K-parallel explanations. But I've just shown that K-ABSTRACT is false, on each of the standard conceptions of abstractness. So this response to the disjunctive threat (response *E*) fails. This point has been obscured, I think, because many inflationists run response *E* without precisely spelling out what they mean by explanatory abstractness.

(More circumspectly: response *E* will fail unless one can find a novel conception of abstractness, a conception that is distinct from the four conceptions I considered in Sect. 5. Haug (2011a, 2011b) promises to provide such a novel conception. Clarke (forthcoming), however, argues that Haug's notion of abstractness fails to be sufficiently different from abstractness as cognitive transparency, despite what Haug hopes.)

7 Conclusion

The aim of this paper has been to clarify the debate over 'in principle' explanatory irreplaceability. Section 1 showed how to ensure that this debate doesn't collapse into a trivial dispute over the definition of an 'idealized agent'. Namely, the debate is most fruitfully read as one about metaphysics: inflationists assert, and deflationists deny, that there are genuine higher patterns or laws. Section 1 distinguished inflationism from other forms of 'anti reductionism' about higher levels, with which inflationism is often conflated. Section 1 also spelled out the 'abstractness argument' for inflationism, as thoroughly and precisely as possible. I noted its inductive form, its variant

³³ Indeed, see Clarke (Forthcoming) for criticisms of arguments that draw the opposite conclusion.

formulations, and its relationship to ‘projectability arguments’ for inflationism. Section 2 then set out a prima facie threat to the abstractness argument—the disjunctive threat. It showed that there are exactly five responses *A B C D E* available to neutralize this threat. And response *A* (deny INFINITE) would trivialize the debate, and undermine the abstractness argument for inflationism, I argued. And Sect. 6 showed that response *E* to the disjunctive threat fails. Response *E* is endorsed by Block, Fodor, Kincaid, Kornblith, Levine, Marras, Pereboom, Putnam, Weslake and at one stage Antony. It claims that higher explanations abstract away from some of the details of their lower *K*-parallel explanations. But this claim is false, I argued, at least on the four standard conceptions of explanatory abstractness outlined in Sect. 5. Although Antony (1999), Clapp (2001) and others have made a similar criticism of response *E*, many philosophers try to resist it, as Sect. 6 showed.³⁴ I hope that the clarificatory work of this paper—especially concerning the four conceptions of abstractness outlines in Sect. 5—allows one to press this criticism as thoroughly, clearly, and compellingly as possible.

In sum, responses *A* and *E* to the disjunctive threat fail, and response *B* (deny NECESS) is relatively unpopular amongst inflationists. This leaves response *C* (deny COMPOUND and SPLIT) and *D* (deny STATES). But Sects. 3 and 4 showed that the availability of these responses turns on how one defines lower states and lower explanations. Therefore, the abstractness argument’s prospects are much brighter for some definitions of lower states and lower explanations than for others. And, although some philosophers are clear about what they mean by the ‘is lower than’ relation, many are not. Indeed, some philosophers often fail to appreciate that the philosophers with whom they engage have a different conception of the ‘is lower than’ relation from their own (Sect. 3). Progress can only be made, therefore, once one accepts that the abstractness argument is really an argument schema, not a single argument.

Acknowledgements I am indebted to Christopher Cowie, Tim Lewens and Nick Shea for their helpful comments on an ancestor of this manuscript; and also to two anonymous referees for their generous suggestions. 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. (2003). Who’s afraid of disjunctive properties? *Philosophical Issues*, 13, 1–21.
- Antony, L. M. (2008). Multiple-realization: Keeping it real. In J. Hohwy & J. Kallestrup (Eds.), *Being reduced: New essays on reduction, explanation and causation* (pp. 164–175). Oxford: Oxford University Press.
- Antony, L. M., & Levine, J. (1997). Reduction with autonomy. *Philosophical Perspectives*, 11, 83–105.
- Bacon, J. (1986). Supervenience, necessary coextensions, and reducibility. *Philosophical Studies*, 49, 163–176.
- Baker, L. R. (1993). Metaphysics and mental causation. In J. Heil & A. R. Mele (Eds.), *Mental causation* (pp. 75–96). Oxford: Oxford University Press.

³⁴ See also Haug’s (2011a) criticisms of response *E* with abstractness understood as the generality of the particular factors an explanation invokes.

- Baker, L. R. (2007). *The metaphysics of everyday life: An essay in practical realism*. Cambridge: Cambridge University Press.
- Bedau, M. A. (2010). 3 Weak emergence and context-sensitive reduction. In A. Corradini & T. O'Connor (Eds.), *Emergence in science and philosophy* (pp. 6–46). London: Routledge.
- Beebe, H., & Sabbarton-Leary, N. (2010). *The semantics and metaphysics of natural kinds*. London: Routledge.
- 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.
- Block, N. (1997). Anti-reductionism slaps back. *Philosophical Perspectives*, 11, 107–132. Originally entitled anti-reductionism strikes back.
- Cartwright, N. (1983). *How the laws of physics lie*. Oxford: Oxford University Press.
- Chalmers, D. (2006). Strong and weak emergence. In P. Clayton & P. Davies (Eds.), *The re-emergence of emergence*. Oxford: Oxford University Press.
- Clapp, L. (2001). Disjunctive properties: Multiple realizations. *The Journal of Philosophy*, 98, 111–136.
- Clarke, C. (Forthcoming). The explanatory virtue of abstracting away from idiosyncratic and messy detail. *Philosophical Studies*.
- Correia, F., & Schnieder, B. (2012). *Metaphysical grounding: Understanding the structure of reality*. Cambridge: Cambridge University Press.
- Elster, J. (1989). *Nuts and bolts for the social sciences*. Cambridge: Cambridge University Press.
- 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*.
- Garfinkel, A. (1981). *Forms of explanation*. New Haven, CT: Yale University Press.
- 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.
- Heil, J. (2003). Levels of reality. *Ratio*, 16, 205–221.
- Jackson, F., & Pettit, P. (1992). In defence of explanatory ecumenism. *Economics and Philosophy*, 8, 1–21.
- Kim, J. (1984). Concepts of supervenience. *Philosophy and Phenomenological Research*, 45, 153–176.
- Kim, J. (1991). Supervenience as a philosophical concept. *Metaphilosophy*, 21, 1–27. Cited as reprinted in Kim (1993, 131–60).
- Kim, J. (1993). *Supervenience and mind*. Cambridge: Cambridge University Press.
- Kim, J. (1998). *Mind in a physical world*. Cambridge: MIT Press.
- 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. (1984). 1953 and All that. A tale of two sciences. *Philosophical Review*, 93, 335–373.
- Lange, M. (2004). The autonomy of functional biology: A reply to Rosenberg. *Biology and Philosophy*, 19, 93–109.
- LaPorte, J. (2004). *Natural kinds and conceptual change*. Cambridge: Cambridge University Press.
- Lewis, D. K. (1972). Psychophysical and theoretical identifications. *Australasian Journal of Philosophy*, 50, 249–258.

- Lewis, D. K. (1994). Reduction of mind. In S. Guttenplan (Ed.), *A companion to philosophy of mind* (pp. 412–431). Oxford: Blackwell. Cited as reprinted in Lewis (1999, 291–324).
- Lewis, D. K. (1999). *Papers in metaphysics and epistemology* (Vol. 2). Cambridge: Cambridge University Press.
- Lipton, P. (1991). *Inference to the best explanation*. London: Routledge.
- List, C., & Menzies, P. (2009). Non-reductive physicalism and the limits of the exclusion principle. *Journal of Philosophy*, 106, 475–502.
- MacDonald, G. (1985). Modified methodological individualism. *Proceedings of the Aristotelian Society*, 86, 199–211.
- 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.
- McGinn, C. (1991). *The problem of consciousness: Essays toward a resolution*. Oxford: Blackwell.
- Mumford, S. (1998). *Dispositions*. Oxford: Oxford University Press.
- Noordhof, P. (2010). Emergent causation and property causation. In C. Macdonald & G. Macdonald (Eds.), *Emergence in mind*. Oxford: Oup.
- Oppenheim, P., & Putnam, H. (1958). The unity of science as a working hypothesis. In H. Feigl, M. Scriven, & G. Maxwell (Eds.), *Concepts, theories and the mind-body problem* (pp. 3–36), Volume 2 of Minnesota studies in the philosophy of science Minnesota: University of Minnesota Press.
- 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.
- Potochnik, A., & McGill, B. (2012). The limitations of hierarchical organization. *Philosophy of Science*, 79, 120–140.
- Prior, E., Pargetter, R., & Jackson, F. (1982). Three theses about dispositions. *American Philosophical Quarterly*, 19(3), 251–257.
- Putnam, H. (1967). Psychological predicates. In W. Capitan & D. Merrill (Eds.), *Art, mind and religion*. Pittsburgh: University of Pittsburgh Press. Cited as reprinted in Putnam (1975a).
- Putnam, H. (1973). Reductionism and the nature of psychology. *Cognition*, 2, 131–146. Cited as revised in Putnam (1975b).
- Putnam, H. (1975a). The nature of mental states. In H. Putnam (Ed.), *Mind language and reality: Philosophical papers* (Vol. 2). Cambridge: Cambridge University Press.
- Putnam, H. (1975b). Philosophy and our mental life. In H. Putnam (Ed.), *Mind language and reality: Philosophical papers* (Vol. 2, pp. 291–303). Cambridge: Cambridge University Press.
- Rosenberg, A. (2001). How is biological explanation possible? *British Journal for the Philosophy of Science*, 52, 735–760.
- Sawyer, R. K. (2002). Nonreductive individualism: Part I-supervenience and wild disjunction. *Philosophy of the Social Sciences*, 32, 537–559.
- Shapiro, L. A. (2000). Multiple realizations. *Journal of Philosophy*, 97, 635–654.
- Shoemaker, S. (2007). *Physical realization*. Oxford: Oxford University Press.
- Silberstein, M. (2012). Emergence and reduction in context: Philosophy of science and/or analytic metaphysics. *Metascience*, 21(3), 627–642.
- Sober, E. (1984). *The nature of selection*. Chicago, IL: Chicago University Press.
- 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.
- Sturgeon, N. L. (1985). Moral explanations. In D. Copp & D. Zimmerman (Eds.), *Morality, reason and truth* (pp. 49–78). Totowa, NJ: Rowman & Allanheld.
- Teller, P. (1984). Comments on Kim's paper. *Southern Journal of Philosophy*, 22, 57–61.
- Van Cleve, J. (1990). Mind-dust or magic? Panpsychism versus emergence. *Philosophical Perspectives*, 4, 215–226.
- Waters, C. K. (1990). Why the antireductionist consensus won't survive the case of classical Mendelian genetics. *Proceedings of the Biennial Meeting of the Philosophy of Science Association*, 1, 125–139.
- Weatherston, B., & Marshall, D. (2014). Intrinsic vs. extrinsic properties. In E. N. Zalta (Ed.), *The stanford encyclopedia of philosophy*. <http://plato.stanford.edu/archives/fall2014/entries/intrinsic-extrinsic/>.
- Weslake, B. (2010). Explanatory depth. *Philosophy of Science*, 77, 273–294.

- Williamson, T. (2007). *The philosophy of philosophy*. Oxford: Blackwell.
- Woodward, J. (2003). *Making things happen: A theory of causal explanation*. Oxford: Oxford University Press.
- Wright, C. (1992). *Truth and objectivity*. Cambridge, MA: Harvard University Press.
- Ylikoski, P. (2011). Social mechanisms and explanatory relevance. In P. Demeulenaere (Ed.), *Analytical sociology and social mechanisms*. Cambridge: Cambridge University Press.
- Zahle, J. (2003). The individualism-holism debate on intertheoretic reduction and the argument from multiple realization. *Philosophy of the Social Sciences*, 33, 77–99.