

The myth of the myth of supervenience¹

David Mark Kovacs

Penultimate draft; please cite the final version, which is forthcoming in *Philosophical Studies*

Abstract: Supervenience is necessary co-variation between two sets of entities (properties, facts, objects, etc.). In the good old days, supervenience was considered a useful philosophical tool with a wide range of applications in the philosophy of mind, metaethics, epistemology, and elsewhere. In recent years, however, supervenience has fallen out of favor, giving place to grounding, realization, and other, more metaphysically “meaty”, notions. The emerging consensus is that there are principled reasons for which explanatory theses cannot be captured in terms of supervenience, or as the slogan goes: “Supervenience Is Nonexplanatory” (SIN). While SIN is widely endorsed, it is far from clear what it amounts to and why we should believe it. In this paper, I will distinguish various theses that could be meant by it, and will argue that none of them is both interesting and plausible: on some interpretations of ‘explanatory’, we have no reason to believe that supervenience is unexplanatory, while on other interpretations, supervenience is indeed unexplanatory, but widely accepted textbook cases of explanatory relations come out as unexplanatory, too. This result raises doubts as to whether there is any interesting sense in which SIN is true, and suggests that the contemporary consensus about supervenience is mistaken.

1. The rise and fall of supervenience

Supervenience is necessary co-variation between two sets of entities (properties, facts, objects, or what have you). To use a very generic definition that doesn’t decide between the various more fine-grained notions of supervenience, we could say that the A-entities supervene on the B-entities just in case there cannot be a difference in the A-entities without some difference in the B-entities (I will give a more precise characterization later on). In the

¹ Acknowledgements For helpful comments on and conversations about earlier versions of this paper I thank Karen Bennett, Nina Emery, Ghislain Guigon, Terry Horgan, Jon Litland, Mike Raven, Oron Shagrir, Alex Skiles, Elanor Taylor, Jessica Wilson, three anonymous referees, and audiences at the 2017 Central APA in Kansas City, the Third Annual Conference of the Society for the Metaphysics of Science at Fordham University, and talks given at the Ben Gurion University of the Negev, the University of Southampton, and the Research Group for the History and Philosophy of Science (RCH HAS) in Budapest.

good old days, supervenience was considered a useful philosophical tool with a wide range of applications. It was most popular in the philosophy of mind² but also enjoyed wide prominence in metaethics³ and epistemology.⁴ Of course, there were lively disputes about exactly which notion of supervenience was most up to task for which purpose. The technical literature in this area was mostly concerned with different forms of property supervenience: strong vs. weak, global vs. local⁵, and various intermediate theses both with respect to pattern of distribution⁶ and strength⁷.

Those days are over. The philosophical landscape has changed dramatically in the last two decades, and supervenience largely fell out of favor, giving place to more metaphysically “meaty” notions such as realization, micro-based determination, and most recently, grounding. The emerging consensus has been that supervenience couldn’t be used to formulate explanatory theses. The following passage is due to Jaegwon Kim, a former champion of supervenience:

“[S]upervenience itself is not an explanatory relation. It is not a “deep” metaphysical relation; rather, it is a “surface” relation that reports a pattern of property variation, suggesting the presence of an interesting dependency relation that might explain it.” (Kim 1988/1993: 167)

² See Davidson 1970/1980, Horgan 1982, Lewis 1983, Chalmers 1996, Jackson 1998, and many, many others.

³ Hare 1952, Klagge 1988

⁴ Van Cleve 1985, Sosa 1991, Conee and Feldman 2004

⁵ See Petrie 1987, McLaughlin 1995, and Sider 1999, to just name a few pieces from the vast literature on the various entailment connections between different forms of supervenience.

⁶ Horgan 1982, Hoffmann and Newen 2007

⁷ Shagrir 2002, Bennett 2004

Concurring, Thomas Grimes (1988) refers to the idea that supervenience relations have explanatory power as the “myth of supervenience”. In the years to come, similar remarks began to appear with increasing frequency⁸, and today they are sprinkled all over the contemporary grounding literature.⁹

Why is this attitude so widespread? Arguably, the first wave of disillusion was gradual and piecemeal: supervenience was used to capture specific and supposedly explanatory theses (such as physicalism), and philosophers came to the conclusion that supervenience was not the right tool to capture these theses. Note, however, that such piecemeal evidence doesn’t support the strong conclusion that supervenience *as such* is non-explanatory; at best, it supports the much weaker claim that it is not a suitable explanatory tool *in such and such cases*. So even if philosophers began to suspect that supervenience was unexplanatory on the basis of a few prominent cases, we would need some general reason to draw the strong conclusion that supervenience is not an explanatory relation.

In subsequent years, a growing number of philosophers have indeed attempted to provide such a reason, by adopting something like the following picture of explanation. There is an elite group of explanatory or determinative relations that “back” or “underlie” explanations. There can be no explanation without the help of one of the elite relations. Causation is one of these relations, and perhaps a few non-causal relations (e.g. micro-

⁸ See DePaul 1987, Horgan 1993, McLaughlin 1995: 18, Shagrir and Hoffman-Kolss 2013, and Endicott 2016: 2207, among others.

⁹ See, for example, Fine 2001, 2012, Schaffer 2009: 364, Raven 2012: 690, deRosset 2013: 2, and Trogon 2013: 101. Even those who attack grounding tend to agree that supervenience has no explanatory value; see Koslicki 2014 and Wilson 2014.

basing, realization, or grounding) are among them, too. Call this the *Backing Model* and its proponents *backing theorists*.¹⁰

Most backing theorists think that supervenience doesn't belong to the elite list of backing relations; it's the wrong *kind* of relation to figure in explanatory theses. Or, as the slogan goes, *Supervenience Is Nonexplanatory* (SIN). Unfortunately it's far from clear what the slogan ultimately comes down to, because philosophers rarely bother to spell out what they mean by the words 'explanatory relation'; the expression is frequently thrown around in the literature as if it was obvious what it meant. But it isn't, and vague appeals to "backing" and "underlying" are too metaphorical and too amorphous to give us a clear answer. Instead of trying to define it, I will undertake a smaller and more manageable task. Whatever SIN exactly amounts to, for it to be a true and interesting thesis there has to be a necessary condition of explanatoriness that supervenience doesn't meet but paradigmatic explanatory relations do. Note that this is a very weak claim. I'm not assuming that 'explanatory relation' has a definition or even an informative set of necessary and jointly sufficient conditions. I'm not even supposing that there is a feature in virtue of which explanatory relations are explanatory. All I'm looking for is a feature whose absence would suffice for a relation to

¹⁰ Backing theorists include Ruben (1990), Kim (1994), Audi (2012), Schaffer (2012, 2016), and many others. Kim (1994) prefers the name 'explanatory realism', which I find misleading (one can be a realist about explanation without thinking of it in terms of underlying relations). Many philosophers use the word 'grounding' for a type of explanation, rather than the underlying relation (e.g. Fine 2012, DeRosset 2013, Dasgupta 2014). Raven (2015) refers to such philosophers as "unionists", and to those who distinguish grounding from metaphysical explanation as "separatists". I should note that while unionists reject the idea that grounding undergirds metaphysical explanation, they don't thereby deny that *something* undergirds it. So, the Backing Model is strictly weaker than unionism. Indeed, even many grounding skeptics, such as Wilson (2014), subscribe to some version of the Backing Model.

not count as explanatory. This is fully consistent with the popular notion that explanatoriness is metaphysically and/or conceptually primitive. In what follows, I will often appeal to the metaphor of the “Great Explanatory Divide”. Past the Divide are all the explanatory relations, and behind it all the non-explanatory ones. I’m happy to give ‘explanatory relation’ to advocates of SIN; all I demand is that whatever they mean by these words, it has to be something that allows us to put supervenience behind the Divide and paradigmatic explanatory relations past it.

The purpose of this paper is to show that accomplishing this task faces formidable difficulties: the most natural interpretations of the words ‘explanatory relation’ either don’t allow us to put supervenience behind the Divide or force us to also put there some paradigmatic explanatory relations. One may draw either of two conclusions from this. The more conservative one is that (for all we know) supervenience *is* an explanatory relation, in whatever sense grounding or causation or realization is. The more radical one is that (barring some hitherto overlooked interpretation of ‘explanatory relation’) the whole idea of the Great Explanatory Divide is misguided: we should abandon the Backing Model altogether and stop thinking of explanation in terms of explanatory relations. Though I’m sympathetic to this more radical option, in the present paper I won’t defend it.

Since SIN is frequently thrown around but rarely spelled out in detail, the question of whether supervenience is an explanatory relation is easily mistaken for other questions in the vicinity. So before discussing SIN itself, I need to introduce some minimal constraints on the notion of supervenience I will work with and distinguish SIN from similar-sounding slogans that fall outside the scope of this paper. I will undertake this ground-clearing task in the next section.

2. Supervenience, explanation, and dependence

Before discussing whether there is an interesting sense in which SIN is true, I should make three clarifications about my target. First, explanation is usually conceived of as a relation between some plurality of facts (the explanantia) and a fact (the explanandum). Hence, I will mostly focus on a special case of fact-supervenience in which some individual fact supervenes on a plurality of facts. Nothing of substance will hang on this, but since this special case of supervenience is formally similar to the explanation relation, focusing on it will considerably streamline the discussion. (More precise definitions could be crafted to capture local, global, regional etc. fact-supervenience. Moreover, we could define supervenience for entities other than facts that can nonetheless stand in intuitively explanatory relations. But these niceties are irrelevant to my goals in the present paper, so I will forgo them).

Second, SIN shouldn't be confused with another thesis that also has to do with explanation: that supervenience connections always stand in need of explanation (perhaps because modal connections in general do).¹¹ The relation between this "no-bruteness" principle and SIN is not entirely clear. For one, in the sections to follow I will consider various interpretations of SIN that don't trade on this principle. For another, it's perfectly consistent to claim that supervenience connections *require* explanation but that they also *provide* a kind of explanation.¹² (The explanans of an ordinary causal explanation, unless it's

¹¹ See Blackburn 1985 and Horgan and Timmons 1992 for influential statements of this idea, and Zangwill 1997 for an opposing view. See McKenzie 2014: 357 for a recent example of someone who mentions SIN and the "no-bruteness" principle in the same breath.

¹² Indeed, Lewis was an ardent denier of brute necessary connections, yet he often seemed to write as if he thought of supervenience as a broadly explanatory notion; see, e.g., Lewis 1983: 358–359.

the Big Bang or some other ultimate first cause, also typically requires further explanation.) In sections 6 and 7 I will consider attempts to infer SIN from the no-bruteness requirement; suffice it to say that the connection is not obvious, and that my main focus in this paper will be SIN.

Third, and finally, SIN is frequently combined with complaints to the effect that supervenience doesn't capture a certain kind of *metaphysical* connection.¹³ These two claims are not independent from each other. Most advocates of SIN either adopt an ontic approach to explanation, according to which explanation is entirely a worldly relation, or a hybrid approach, according to which it's a relation with worldly as well as epistemic/psychological aspects.¹⁴ On either kind of view, focusing on explanation is one possible way of making precise exactly *what* metaphysical connection supervenience ought to (but supposedly doesn't) capture. While there are other options, they don't account as well for the recent decline of supervenience. For example, it would clearly be misguided to complain that supervenience fails to capture *any* metaphysical connection whatsoever. Supervenience is modal co-variation, so any notion of supervenience in which the modality at issue is metaphysical necessity at least guarantees metaphysically necessary co-variation – surely a metaphysical connection if anything is. A more targeted complaint could be that supervenience fails to capture the kind of “nothing-over-and-above”-ness that (for instance) has been thought to be an important desideratum of physicalism.¹⁵ But the words ‘nothing over and above’ are notoriously unclear; one can of course make the phrase clearer by stipulating how one wants to use it, but I doubt there is a pre-theoretical notion in the

¹³ Thanks to an anonymous referee, who drew my attention to this related concern.

¹⁴ I borrow this tripartite taxonomy from Trogdon forthcoming.

¹⁵ [Acknowledgment omitted].

vicinity that is sufficiently well understood to structure the debate over the utility of supervenience around it. The other problem is that although the received view today is that supervenience doesn't guarantee "nothing-over-and-above"-ness (whatever the phrase exactly means)¹⁶, one can think the same about the notions that came to replace it without taking this to be a reason to shun those notions.¹⁷ For this reason, it seems unpromising to construe the debate over the utility of supervenience as one about whether it guarantees "nothing-over-and-above"-ness.

One further metaphysical relation one might think is not captured by supervenience is ontological dependence. This requires a bit more spelling out, since this notion is liable to cause confusion in the present context. Those who object to supervenience on the basis that it's not an explanatory relation also often say that it's not a "dependence relation". Sometimes "dependence relation" is used interchangeably with "explanatory relation", and on this usage, the complaint doesn't differ from SIN. However, this use of the phrase "dependence relation" is at the very least terminologically misleading. Ontological dependence has its own vast literature¹⁸, a glance at which should make it clear that being a "dependence relation", i.e. being a species of ontological dependence strictly understood, is not obviously the same thing as being an explanatory relation.¹⁹ To put things as neutrally as

¹⁶ See Yablo 1992, Horgan 1993, and Wilson 1999, among others.

¹⁷ For instance, deRosset (2010) and Audi (2012) deny that grounding implies "nothing-over-and-above"-ness, and Ridge (2007) is plausibly read as denying that realization does.

¹⁸ See Correia 2008, Koslicki 2013, and Tahko and Lowe 2015 for useful surveys on ontological dependence.

¹⁹ A number of authors have made this point by now: see, for instance, McKenzie forthcoming, Schnieder forthcoming, and Author forthcoming-3. Thomas Grimes (1991) recognized the conceptual distinction between dependence and determination much earlier on, although (as was typical at the time) he tried to capture it with finer-grained notions of supervenience.

possible, explanation and dependence are both relations of priority in some sense of the word ‘priority’, and they both involve a modal connection (a kind of “constraining”) between the prior and the posterior thing. Different species of explanation and dependence vary with respect to the strength and scope of the modal connection they impose, but it is an invariant feature of all these notions across the board that in cases of ontological dependence it’s the posterior thing (the dependent) that imposes such constraints on the prior thing (the dependee), whereas in cases of explanation it is the prior thing (the explanans) that puts constraints on the posterior thing (the explanandum).²⁰ In other words, if the dependent is in a certain way, the dependee has to be in a certain way, and if the explanans is in a certain way, the explanandum also has to be in a certain way.

Whether there is any interesting connection between these two types of priority/constraining relations is a substantive question I cannot discuss here. What I want to emphasize is that most notions of supervenience clearly don’t have the modal profile normally assigned to any species of ontological dependence: supervenience implies that no difference in the supervenient is possible without some difference in the supervenience base, but (unless extra conditions are added, which don’t follow from the notion of supervenience as such) there may well be some difference in the supervenience base without a difference in the supervenient. It is exactly this feature that made supervenience attractive to those who wanted to leave room for multiple realizability: they wanted a relation that given a base would fix the higher-level phenomena, but which would also allow for various possible bases

²⁰ For example, Correia 2008 distinguishes between “rigid dependence”, where the dependee requires the existence of some particular dependent thing, and “generic dependence”, where it only requires a certain *kind* of dependent thing. This is a fairly standard distinction in the ontological dependence literature.

to give rise to the same phenomena.²¹ For this reason, there are many contexts outside the specialized literature on ontological dependence where the word ‘dependence’ is best interpreted as meaning something closer to determination or explanation. For example, philosophers of mind with non-reductive physicalist leanings often claim to be in the business of searching for the relation that captures the “dependence” of mental states on physical states. But in my view, we shouldn’t take this at face value: non-reductive physicalists don’t have to think that whenever there is a mental state, there must be some particular accompanying physical state or even some particular *kind* of physical state. What they need instead is a relation that guarantees that if there is a certain physical state then there is also an accompanying mental state. In short, what they are interested in is a non-causal kind of explanation, rather than dependence, even if (somewhat carelessly, I say) they often use the word ‘dependence’.

Now, as I said, some use the expression ‘dependence relation’ much more broadly than I do, and I don’t want to get involved in a verbal dispute over how to use it. But when used narrowly for ontological dependence and its various species, it seems clear that the reason supervenience fell out of favor has nothing to do with whether it’s a dependence relation in my sense – not least because it’s just too obvious that it isn’t. And while there may well be further interesting questions about what kinds of metaphysical connections supervenience can or cannot capture, SIN has at least been *a* major complaint against supervenience, even if not the only complaint. Accordingly, in the next few sections I will focus on SIN and will survey a number of things one might mean by it: that supervenience doesn’t have the right formal properties to track the directedness of explanations (section 3), that it’s too “coarse-

²¹ See Correia 2005: 36–37 for a similar point about the relation between supervenience and ontological dependence; cf. Grimes 1991.

grained” (section 4), that its presence doesn’t guarantee explanation (section 5), that it cannot be “brute” but always requires some underlying explanatory relation (section 6), and that its explanatory value solely derives from such relations (section 7). While none of these claims is uncontroversial among proponents of the slogan, it’s safe to say that they all take at least one of them to be both true and sufficient for supervenience to be unexplanatory in the sense they are interested in. However, I will argue that once we take a closer look at these interpretations, we find that none of them carves out a sense in which SIN is both true and interesting.

3. The “Wrong Formal Properties” worry

One frequently heard warning about supervenience is that it has the wrong formal properties to qualify as an explanatory relation. Explanatory relations are asymmetric, irreflexive, and non-monotonic, we are told, whereas supervenience is reflexive, non-symmetric, and monotonic. Several authors, especially in the grounding literature, cite these features of supervenience in arguing that it’s not an explanatory relation.²² Call this the “Wrong Formal Properties” worry.

The worry is misplaced. The reason for this is *not* that contrary to popular assumption explanation is neither asymmetric nor irreflexive.²³ I think it is, and at any rate, I don’t want to rest my case on controversial examples of self-explanation. Instead, following Berker (forthcoming), I point out that the problem has an easy fix: we can simply replace supervenience with *proper* supervenience, where fact A properly supervenes on the B-facts iff

²² See, for instance, Schaffer 2009: 364, McLaughlin and Bennett 2011: §3, §5, and Raven 2012: 690.

²³ *Pace* Bliss 2014 and Wilson 2014. See Author forthcoming-1 for a general defense of the irreflexivity of explanation and explanatory relations.

(1) A supervenes on the B-facts, (2) not vice versa, and (3) A doesn't supervene on any proper subset of the B-facts. Clause (2) ensures asymmetry and irreflexivity, and clause (3) ensures non-monotonicity.

A few remarks are in order here. First, clause (3) makes proper supervenience not only non-monotonic but also minimal, i.e. it implies that any addition to an explanatorily sufficient supervenience base doesn't merely fail to guarantee but also "kills" explanation.²⁴ It is controversial whether all explanations are minimal in this sense, and I'm not claiming that *every* explanatory instance of supervenience is a case of proper supervenience.

Second, I'm also not asserting that the formal restrictions I built into proper supervenience *suffice* to make it an explanatory relation. This claim is obviously false, since many relations that have the requisite formal properties (for example, *having smaller net mass than*) are not explanatory. Indeed, I'm not even claiming that the formal properties of proper supervenience give us *any* reason to think that it's explanatory. Rather, I'm looking for a reason to put supervenience behind the Great Explanatory Divide, and I'm pointing out that the "Wrong Formal Properties" worry doesn't generalize to proper supervenience and so doesn't put it behind the Divide. It doesn't follow that proper supervenience is an explanatory relation. What follows is that the "Wrong Formal Properties" worry fails to show that it isn't, or to explain why it isn't even if it isn't.

Third, one might complain that proper supervenience has the right formal properties simply as a matter of definition. We introduced the expression 'proper supervenience' for certain special cases of supervenience, which display the desired formal properties. One might object that due to its stipulative nature, proper supervenience is somehow an arbitrary

²⁴ Salmon (1977/1997: 95–96) and Audi (2012: 699–701) accept the minimality constraint; Fine (2012) and Litland (2017) reject it. [Acknowledgments omitted]

restriction on supervenience. But if this is a worry, it may also apply to paradigmatic explanatory relations like grounding. Many philosophers recognize a distinction between grounding in the narrow, explanatory sense, and a broader notion in the vicinity suitable for certain technical purposes. Moreover, many of them define the narrow, explanatory notion as a restriction on the broader notion, or are at least open to doing so. For example, Schaffer suggests that grounding proper could be defined in terms of a non-irreflexive notion: x properly grounds y iff x (plainly) grounds y and $x \neq y$.²⁵ In a similar fashion, Fine chooses weak ground as his preferred primitive and proposes to define strict ground in terms of it: a set of truths Δ strictly grounds C iff Δ weakly grounds C and C doesn't weakly ground any member of Δ , by itself or with other statements.²⁶

On any of these definitions, the narrow, explanatory notion of grounding has its formal properties by stipulation no less than does proper supervenience. Yet nobody seems to think that this feature makes Fine's notion of strict ground, or Schaffer's notion of proper grounding, unexplanatory. And grounding is not special in this regard. Proper parthood, too, is often considered an explanatory relation.²⁷ But it's common to define proper parthood in terms of parthood (x is a proper part of y iff x is a part of y and $x \neq y$), which makes the asymmetry and irreflexivity of proper parthood a matter of stipulation.

Fourth, perhaps the analogy with weak ground does show that, strictly speaking, supervenience itself is not an explanatory relation. However, that doesn't yet give us an

²⁵ Schaffer 2009: 374

²⁶ See Fine 2012: 52. Fine uses 'grounding' for metaphysical explanation, rather than the underlying relation, but this is irrelevant to the point I'm making.

²⁷ See Kim 1994: 67, Markosian 2005: §3, and Cameron 2014. Wilson (2014: 539) considers the more restricted *causal* composition relation explanatory.

interpretation of SIN on which it's both true and interesting. Surely advocates of SIN don't think that supervenience is unexplanatory only in the same sense in which weak ground is; nor do they suppose that proper supervenience is explanatory in the same sense in which strict ground is.

One might object instead to the artificially imposed non-monotonicity requirement. Explanatory relations, the complaint goes, are non-monotonic because they satisfy a *relevance* constraint. And although proper supervenience is non-monotonic, it's false that each member of a proper supervenience base is relevant to the supervenient. But this objection is unlikely to advance the dialectic. The relevance constraint in question cannot simply be understood in terms of relevant implication, since many paradigmatic explanatory relations violate it: the obtaining of an effect isn't relevantly implied by the obtaining of the totality of its causes.²⁸ So, 'relevance' has to mean something like *explanatory* relevance (as in Guigon forthcoming). But it's unhelpful to be told that explanatory relations are non-monotonic, and that they are non-monotonic because the relata occupying their first argument place are explanatorily relevant to the relatum occupying their second argument place. This might be true, but it doesn't give us an independent necessary condition of explanatoriness that (proper) supervenience violates and paradigmatic explanatory relations satisfy.

None of the foregoing discussion shows that SIN is false or that the formal properties of proper supervenience make it an explanatory relation. I'm only making the weaker claim that the "Wrong Formal Properties" worry is a bad reason to think that it isn't, and therefore it doesn't provide a plausible interpretation of SIN. If supervenience is not an explanatory relation, this is for some reason other than that it has the wrong formal properties.

²⁸ Relevant implication is not only unnecessary but also insufficient for explanation; see Salmon 1977/1997: 95.

4. The “Coarse-grainedness” worry

Perhaps the main concern with supervenience is that it ignores distinctions that genuine explanatory relations respect. One common way of putting this is that supervenience is an intensional relation, whereas explanatory relations are hyperintensional.²⁹ That is, whenever supervenience holds between certain relata, it also holds between pairwise necessarily coexisting relata. Not so for grounding. For example, according to most grounding theorists, P&Q both supervenes on and is grounded in P and Q taken together. Moreover, P and Q taken together obtain just in case $\sim\sim P$ and $\sim\sim Q$ do. Yet P&Q supervenes on but isn't grounded in $\sim\sim P$ and $\sim\sim Q$ taken together.

However, we have reason to doubt that the intensionality of supervenience is really what is at issue in the coarse-grainedness worry. This is because even if supervenience itself is intensional, proper supervenience (as defined in the previous section) isn't. Suppose fact A properly supervenes on the B-facts, A* obtains just in case A does, and the B*-facts obtain just in case the B-facts do. This doesn't guarantee that A* properly supervenes on the B*-facts. Suppose, for instance, that A and B are contingent facts, B* is a conjunction of B and the fact that $2+2=4$, and A* is the conjunction of A and the fact that there are no round squares. Since A* and A, and the B*-facts and the B-facts, are pairwise cointensional, A's (proper) supervenience on the B-facts guarantees A*'s supervenience on the B*-facts. Yet A* doesn't properly supervene on the B*-facts, since the latter have a proper subset (the B-

²⁹ Perhaps causation is an exception, and only non-causal explanatory relations are hyperintensional; this doesn't matter for the argument. For versions of the “Coarse-grainedness” worry, see Schaffer 2009: 364, 2016: 52, Raven 2015: 325, Bliss and Trogon 2014: §4, and McKenzie 2014, forthcoming.

facts) on which it also supervenes. Therefore proper supervenience is hyperintensional, after all.

Importantly, the point is not that if proper supervenience is hyperintensional then it's explanatory. Indeed, proper supervenience itself is a relation that can obtain between a pair of relata without those relata standing in the explanation relation. The example I began with is a case in point: P&Q properly supervenes not only on P and Q taken together, but also on $\sim\sim P$ and $\sim\sim Q$ taken together; yet the dominant view is that $\sim\sim P$ and $\sim\sim Q$ don't jointly explain P&Q.³⁰ So, the point is not that its hyperintensionality makes proper supervenience an explanatory relation, but rather, that blaming the relation's supposed unexplanatoriness on its intensionality is a bad diagnosis. If you have the intuition (as I expect most of my readers to have it) that proper supervenience is also unexplanatory, the reason must lie elsewhere.

The textbook statement of the coarse-grainedness worry is that explanatory relations recognize distinctions that supervenience doesn't. According to the hyperintensionality interpretation, the distinctions in question are those that cut more finely than modal notions do, but we have seen above that this doesn't get to the heart of the worry. Moreover, similar problems are likely to arise for other attempts to make precise exactly which relevant distinctions are the ones that supervenience is insensitive to. After all, we can keep defining more and more fine-grained notions of supervenience by simply incorporating extra conditions by *fiat* (for example, we could introduce a ban on double-negated properties in the supervenience base to put a wedge between P, Q and $\sim\sim P$, $\sim\sim Q$). One might be tempted to respond that these doctored notions are *ad hoc*, but this would miss the mark. "Supervenience" is not a word of ordinary English; *all* types of supervenience are technical

³⁰ Thanks to an anonymous referee for helping me clarify the argument here .

notions invoked and defined to play certain theoretical roles. It is perfectly acceptable to fine-tune such notions until they can perform the tasks we assigned to them. The letter of the coarse-grainedness worry is merely that supervenience fails to recognize certain distinctions we want to draw, and extra stipulations would adequately address *this* worry.

This being said, I think there is a deeper concern behind the coarse-grainedness worry that isn't addressed merely by putting extra restrictions on supervenience. I suspect that any non-trivial restriction (that is, any restriction short of directly stipulating that the supervenience relata stand in the explanation relation) will face confounding cases in which we have supervenience with all the relevant restrictions but intuitively still don't have an explanation. What this seems to suggest is that the "Coarse-grainedness" worry is an imperfect way of getting at a more basic concern, which we can express more directly by simply saying that genuine explanatory relations distinguish between sets of facts that stand in the explanation relation and sets of facts that don't. (This can be seen as a limiting case of the coarse-grainedness worry: there are indeed some distinctions that explanatory relations capture and supervenience doesn't, but there is no more basic way of saying what these distinctions are other than that they are the ones that track explanation.)

While this might be a serious worry, it is misleading to describe it as a problem with coarse-grainedness *per se*; the real issue it expresses is that supervenience doesn't guarantee explanation. I turn to this objection in the next section.

5. The "No Guarantee of Explanation" worry

As I mentioned in section 1, according to the popular Backing Model explanatory relations "back" explanations. While it's rarely spelled out what this means, one might plausibly think that it requires at least the following:

(*Guarantee Condition*, GC) If R is an explanatory relation, then necessarily, for any $x_1 \dots x_n$ and any y , if $R(x_1 \dots x_n, y)$ then $x_1 \dots x_n$ explains y .³¹

Given GC, one might think that supervenience cannot be an explanatory relation, since there are well-known cases of supervenience without explanation. For example, any mathematical fact globally supervenes on the mental facts: it couldn't be different without the mental facts being different, since as a mathematical fact it's necessary and so couldn't be different *period*. Likewise, any arbitrary fact f_a supervenes on any set of facts one of which is the conjunction $f_1 \& f_2 \& \dots \& f_n$, where $a \in \{1, \dots, n\}$. Yet it seems obvious that conjunctive facts don't explain their conjuncts. Examples could be multiplied; there are obviously plenty of cases of supervenience without explanation.³²

The problem with GC is that paradigmatic explanatory relations violate it. Take, for instance, causation. It's not too hard to think of cases in which an event is causally relevant to another event without being explanatorily relevant to it. Some causes are explanatorily irrelevant because they only influence fine-grained features of the explanandum event far removed from our explanatory concerns.³³ For example, a complex series of events led to Rasputin's death: he was served poisoned teacakes, shot twice, and finally thrown into the Neva river, where he drowned. The primary cause of his death was the drowning, though the

³¹ GC could be fine-tuned to accommodate intuitively explanatory relations between entities that cannot themselves stand in the explanation relation, but since I chose to focus on fact-supervenience, I will forgo this task here.

³² Cf. Audi (2012b: 687–688), Schaffer (2009: 364–365, 2016: f52), Rosen (2010: 110–114), and Fine (2012: 41).

³³ Ruben 1990: Ch. 5, Ch. 7: 187–193

poisoning and the bullet wounds may have hastened the process. But while the gravitational influence of Mars and the day's pollen count exerted causal influence on Rasputin's death, they were explanatorily irrelevant to it.³⁴ Another case in which causes fail to explain is when they are too far removed in the causal chain from the explanandum event. For example, even in a deterministic world, citing the Big Bang would not make for an explanation of why Jimmy was late from school on a certain day, even though the Big Bang was clearly in the chain of causes that led to his lateness, and perhaps even fully determined it.³⁵

A natural response to these counterexamples is to play down the causal relevance of non-explanantia. For example, one might object that once we accept that the gravitational force of Mars was explanatorily irrelevant to Rasputin's death, or that the Big Bang was explanatorily irrelevant to Jimmy's lateness from school, we should start feeling some pressure to also say that these events were causally irrelevant to the respective explananda. However, in my view, if this is so it rather points to a significant problem for the Backing Model. As Strevens (2008: Chs. 2, 6) observes, in ordinary discourse we often don't clearly distinguish between causation and causal explanation: in most everyday contexts, an utterance of the form 'c causes e' is really just an elliptic way of saying that c provides a causal explanation for e. If there is a natural physical relation that deserves to be called 'causation', we shouldn't expect that relation to line up with our intuitions about the classic puzzle cases about causation. But then, the pressure we might feel to deny the causal relevance of any explanatorily irrelevant event likely stems from our tendency to confuse causation with causal explanation.

³⁴ See Strevens 2008: 46 for this example.

³⁵ See Lipton 2001: 49. See also Author forthcoming-2 for more on these cases.

This tendency is by no means confined to causation. Take, for instance, realization, which was probably the most widely discussed non-causal explanatory relation before grounding appeared on the scene. While grounding is usually taken to resist analysis, the history of realization is marked by a series of increasingly complicated definitions. I think these definitions were often motivated by implicitly explanatory considerations. Sydney Shoemaker's recent work is a case in point. He first gives the following simple gloss on realization: property P realizes property Q just in case (1) the forward-looking causal powers of Q are a proper subset of the forward-looking causal powers of P, whereas (2) the backward-looking causal powers of P are a proper subset of the backward-looking causal powers of Q.³⁶ But then, he notes that this first gloss would imply that every conjunctive property is a realizer of each of its conjuncts. "Obviously this must be avoided", he says, and goes on to propose a more complicated definition that avoids this consequence. However, we should ask *why* it is so important to avoid it. And it's hard to think of an answer other than that the original, simpler definition doesn't fit our explanatory intuitions.

If this is the answer, it's unsurprising that realization as Shoemaker defines it guarantees explanation: the notion is tailor-made to our explanatory concerns and may more accurately be labeled the *realization-explanation* relation. Similarly, if certain accounts of causation seem to imply that causation is always sufficient for causal explanation, we have reason to think that those accounts confuse the worldly, natural relation of causation with the part-metaphysical, part-epistemic concept of causal explanation. And while I have no space to

³⁶ Shoemaker 2007: 13ff. This is one version of the "subset account of realization", which has first been formulated by Wilson (1999). I use Shoemaker's version rather than Wilson's because it better illustrates the kind of conflation I'm talking about.

argue for this here, I suspect that the situation is similar with other putative explanatory relations.³⁷

At the beginning of this paper, I complained about the obscurity of the expression “explanatory relation”. The past sections should already have gone a long way toward substantiating this complaint: that SIN has so many salient (and actually defended) interpretations indicates that the phrase lacks a clear, universally agreed-upon meaning. But perhaps this also has an advantage: it might be possible to interpret the notion of “backing” in a way that avoids commitment to GC. That is, perhaps we shouldn’t demand that for a relation R to back explanations and count as explanatory, any x s that bear R to some y also

³⁷ I think that a similar difficulty besets a strategy that is the mirror image of the one I considered above: that of playing up the explanatory relevance of causes. This is the route chosen by Skow (2016a), who rejects putative counterexamples to GC by appealing to the pragmatics of explanation. The basic idea is that explanations are answers to ‘Why?’-questions; therefore, a counterexample would be a case in which citing a cause of e in response to ‘Why did e occur?’ would fail to answer the question. Skow argues that all instances of this can be explained away as cases where citing a cause does provide an answer, just not the one we were looking for. However, I doubt this strategy succeeds, since I think that both ‘Why?’-questions and the ‘because’-statements that answer them are ambiguous between causation and causal explanation (see Strawson 1985). As a result, while there is a reading on which ‘ c caused e ’ implies ‘ e occurred because c occurred’, and also a reading on which ‘ e occurred because c occurred’ implies ‘the occurrence of c explains the occurrence of e ’, I deny that there is any reading on which ‘ c caused e ’ automatically implies ‘the occurrence of c explains the occurrence of e ’. It’s worth noting that in the background there is a deeper disagreement between Skow and me. Elsewhere, Skow (2016b: Ch. 1) argues that instead of theorizing about explanation, philosophers should have been focusing on ‘Why?’-questions from the get-go. I disagree; I take explanation to be the better understood of the two notions, which makes me skeptical of calls to restructure our inquiries around ‘Why?’-questions.

explain y .³⁸ Indeed, if the argument in this section was convincing thus far, it's doubtful that any interesting metaphysical relation satisfies this criterion.

Now, we can of course abandon GC and still adopt a version of the Backing Model. However, it's difficult to see how such a weakened version could allow us to give any argument in the vicinity of the "No Guarantee of Explanation" worry. That worry essentially trades on the idea that supervenience cannot be an explanatory relation because there are straightforward *counterexamples* to the putative link between supervenience and explanation. But "counterexample"-talk makes sense only against the backdrop of a modal link between supervenience and explanation that we require in order for the former to count as explanatory. Without such a link, it is difficult to make sense of the dialectical role of problem cases in the literature. After all, there are plenty of well-behaved cases of supervenience with all the trappings of genuine explanations: the supervenience base concerns a more basic part of reality; the relation between it and that which supervenes is asymmetric; it is intuitive that the supervenience base explains the supervenient and has no redundancy in it; and so on. That there are so many well-behaved instances of supervenience that *do* track our explanatory concerns is, I take it, the main reason that decades were spent on cooking up increasingly fine-grained and counterexample-resistant notions of supervenience before the project was finally abandoned and SIN has become metaphysical orthodoxy.

³⁸ This possibility was suggested to me by an anonymous referee.

To be sure, just because there is often an explanatory connection between the subvenient and the supervenient, it doesn't follow that supervenience is itself an explanatory relation. Nor did I say that it did. The point is that GC at least offers independent means to settle this question, which would have allowed us to sidestep the difficult issue of what exactly makes a relation explanatory. By allowing that a relation can back explanations despite not always guaranteeing explanation, we are left without an independent test against which we could evaluate the hypothesis that supervenience is an explanatory relation.

Here's where the foregoing discussion leaves us. The "No Guarantee of Explanation" worry contends that since there can be supervenience without explanation, supervenience is not an explanatory relation. I tried to show that GC, which the argument relies on, classifies some paradigmatic explanatory relations as nonexplanatory, too, and therefore fails to carve out a sense in which SIN is both true and interesting. I then considered the question of whether we could accept the Backing Model but reject GC. I argued that while there is nothing incoherent in doing so, the resulting view would leave us without means to evaluate putative counterexamples to the supposed explanatoriness of supervenience. If a relation's being explanatory doesn't imply a necessary connection between that relation and explanation, then the claim that supervenience is an explanatory relation is not the sort of thing that could be refuted by counterexamples.

This result encourages us to look for alternative interpretations of SIN that don't rely on intuitive "counterexamples" to the supervenience-explanation link. In the next two sections, I will consider two such interpretations. According to each, supervenience is a "superficial" relation in a sense in which genuine explanatory relations aren't, and this is what puts those relations and supervenience on different sides of the Great Explanatory Divide.

6. The “Surface Relation” worry

According to a commonly heard complaint, supervenience isn't explanatory because it's a “surface relation”: whenever it *seems* explanatory, some underlying relation ensures supervenience and does the real explanatory work.³⁹ For example, wholes have their properties because of the properties of and relations among their parts; moreover, their intrinsic properties supervene on the intrinsic properties of and relations among their parts. Yet it's not for this supervenience connection that the parts' properties and relations explain the whole's properties. Instead, the whole's intrinsic properties are *microbased*, i.e. identical to the property of having parts with such and such properties, standing in such and such relations. This is what explains both the supervenience and the explanation relation between a whole's having some properties and its parts having certain properties and standing in certain relations.⁴⁰ Likewise for other “underlying” relations that could account for supervenience, for example realization or constitution: when supervenience comes with these relations, they do the real explanatory work.⁴¹

It's one thing to say that whenever supervenience is accompanied by explanation, it's undergirded by some other relation. It's another thing to conclude that supervenience is therefore unexplanatory. To succeed at putting supervenience and paradigmatic explanatory relations on different sides of the Great Divide, this inference would need to rely on a general principle: for *any* relation, if every instantiation of that relation is explained by the

³⁹ See Kim 1988/1993: 167 and Horgan 1993.

⁴⁰ This example is taken from Kim 1998: 84ff.

⁴¹ See also Koslicki 2004: 336 n13–14 and Melnyk 2003: Chs. 1–2, respectively, for similar remarks on constitution and realization.

instantiation of some underlying explanatory relation ('explanatory' in whatever sense paradigmatic explanatory relations are), then the first relation isn't explanatory in that sense. Without this general principle it would be wholly arbitrary to single out supervenience as unexplanatory.

But there is a problem with this general principle: it categorizes too many relations as unexplanatory. Most accounts of realization define the realization relation, instead of treating it as a primitive. It's natural to read these accounts as also saying that every instance of realization is explained by the pattern of relations we used to define it. This seems plausible in the case of the Wilson-Shoemaker account, but it's even clearer on views that define realization in explicitly mereological terms.⁴² Yet adherents of SIN would hardly think that just because realization can be explained further, this is a reason to deem it unexplanatory.

Similar remarks apply to programming (Jackson and Pettit 1990), which is usually understood in terms of micro-basing: a property P programs for another property Q just in case it's nomologically necessary that whenever P is instantiated, so are some properties $P_1 \dots P_k$ that micro-base P, and which under some ceteris paribus conditions cause instantiations $Q_1 \dots Q_n$ that in turn micro-base Q.⁴³ Again, programming is generally considered an explanatory relation (the one that accounts for program explanations) despite

⁴² Ehring (2011: Ch. 5), for instance, understands realization in terms of the subclass relation (which according to him is a special case of parthood) between classes of tropes. He even explicitly offers this account as a "*metaphysical explanation* for why the sets of causal powers of mental properties stand in the subset relation to the sets of causal powers of certain physical properties—the core of the realization relation, according to the Subset Account" (172, emphasis in the original).

⁴³ I borrow this characterization from [Acknowledgment omitted, p.c.]; Jackson and Pettit don't themselves provide a precise definition of programming.

the fact that each instance of programming is explainable by the instantiation of a complex pattern of causal and micro-basing relations.

One final example: some philosophers posit an intimate connection between grounding and essence; whenever $f_1 \dots f_n$ ground g , this is due to the fact that it lies in g 's essence that if $f_1 \dots f_n$, then g .⁴⁴ However, according to Fine (2012), essence itself corresponds to a certain kind of explanation: whereas grounding provides explanation-‘How?’, essence corresponds to explanation-‘What?’. If we put these views together, we get that grounding is always undergirded by another explanatory relation, essential connectedness. Yet, while not all grounding theorists are grounding essentialists, they all seem to agree that the explanatory status of grounding isn’t hostage to the falsity of this family of views.

As in the previous sections, these examples are not meant to demonstrate that supervenience is an explanatory relation. They are meant to show only that the “Surface Relation Worry” is unlikely to classify supervenience as a non-explanatory relation without also classifying some paradigmatic explanatory relations as non-explanatory, and thereby fails to give justice to SIN.

7. The “Superficial Explanation” worry

One might complain that I didn’t present the “Surface Relation” worry in its strongest form. The problem with supervenience is not just that whenever it seems explanatory, some deeper explanatory relation backs it; on the top of that, it entirely owes its explanatory value to that of the underlying relations. Call this the “Superficial Explanation” worry.⁴⁵ Note that to evaluate this worry, we don’t need to say whether superficially explanatory relations are

⁴⁴ See Rosen 2010, Audi 2012 and Dasgupta 2014 for similar views; cf. Fine 2012 and Correia 2013.

⁴⁵ [Acknowledgment omitted].

“really” explanatory. Throughout the paper, my strategy has been to give the (to my mind obscure) expression “explanatory relation” to my opponent and argue that whatever it means, SIN is either false or uninteresting. Accordingly, below I will argue that whatever feature is supposed to make supervenience only superficially explanatory, it is either a feature that supervenience lacks or one it shares with some paradigmatic explanatory relations.

On the face of it, the “Superficial Explanation” looks more sophisticated than the “Surface Relation” worry. But I think this is mainly due to the obscurity of the expression ‘owes its explanatory value to’. What does it mean for a relation to owe its explanatory value to some other relation? Here is a first stab:

(Superficial-1) A relation, R , is superficially explanatory iff for any $x_1 \dots x_k, y$, if $[R(x_1 \dots x_k, y)]$ backs $[x_1 \dots x_k \text{ explain } y]$, there is some deeper explanatory relation, D , and some $w_1 \dots w_n$ and z , such that $[D(w_1 \dots w_n, z)]$ backs $[x_1 \dots x_k \text{ explain } y]$

On this interpretation, the “explanatory value” of a relation simply refers to the explanations that relation “backs”. So, supervenience is superficially explanatory just in case the explanations it backs are also backed by facts about some deeper explanatory relation. Unfortunately, this proposal brings no real improvement over the “Surface Relation” worry. To see this, we should first make more precise the ‘backing’ (‘underlying’, ‘undergirding’, etc.) talk so prevalent in the grounding/explanation literature. I have postponed this until now, but to make sense of the worry we do need a clearer grasp of what it means for an explanatory relation to back an explanation. The most plausible interpretation I can think of is that ‘backing’ is simply another word for (non-causal) explanation: $[R(w_1 \dots w_n, z)]$ backs

$[x_1 \dots x_k \text{ explain } y]$ just in case $[R(w_1 \dots w_n, z)]$ explains $[x_1 \dots x_k \text{ explain } y]$.⁴⁶ This turns Superficial-1 into

(*Superficial-2*) A relation, R , is superficially explanatory iff for any $x_1 \dots x_k, y$, if $[R(x_1 \dots x_k, y)]$ explains $[x_1 \dots x_k \text{ explain } y]$, there is some deeper explanatory relation, D , and some $w_1 \dots w_n$ and z , such that $[D(w_1 \dots w_n, z)]$ explains $[x_1 \dots x_k \text{ explain } y]$

The problem with Superficial-2 is the following. We have seen in the previous section that facts about some paradigmatic explanatory relations are often thought to be backed by (that is, explained by facts involving) other explanatory relations: realization might be explained by parthood, programming by micro-basing, and perhaps grounding by essential connectedness. But non-causal explanation is plausibly transitive.⁴⁷ So, the instantiation of whichever deep explanatory relation explains the instantiation of one of these paradigmatic explanatory relations also explains the explanations those relations supposedly back. For example, if grounding explains metaphysical explanation and essential connectedness explains grounding, then essential connectedness also explains metaphysical explanation. Thus, grounding (if “backed” by essential connectedness) comes out as only superficially explanatory on this reading. More generally: on Superficial-2, any paradigmatic explanatory relation backed by some other explanatory relation comes out as only superficially explanatory. I take this to be a bad result.

⁴⁶ See Schnieder 2010, 2015 and Author forthcoming-2 for this interpretation of ‘backing’-talk.

⁴⁷ Some deny that grounding is transitive, and presumably they would deny the transitivity of non-causal explanation, too (Schaffer 2012). But even if they are right, it’s fairly plausible that the particular cases I’m discussing here are transitive.

The above line of reasoning generalizes: for any explanatory relation R1 backed by some other explanatory relation, R2, the explanations the R1-facts explain are also explained by the R2-facts. Thus, any paradigmatic explanatory relation backed by some other explanatory relation is superficial according to Superficial-2 and, given the “backing = explanation” interpretation, according to Superficial-1 as well. This is a bad result for proponents of SIN, since as we have already seen, the “Surface Relation” worry fails to put supervenience and all paradigmatic explanatory relations on different sides of the Great Explanatory Divide. But given the reasoning above, Superficial-2 will replicate this result.

Perhaps, however, this is not the strongest construal of the “Superficial Explanation” worry. Perhaps by “the explanatory value” of a putative explanatory relation we should mean not the explanations it backs, but instead the facts to the effect *that* the relation backs such and such explanations. That is,

(*Superficial-3*) A relation, R, is only superficially explanatory iff for any $x_1 \dots x_k, y$, if $[R(x_1 \dots x_k, y)]$ explains $[x_1 \dots x_k \text{ explain } y]$, then there is some deeper explanatory relation, D, distinct from R, and some $w_1 \dots w_n$ and z , such that $[D(w_1 \dots w_n, z)]$ explains $[[R(x_1 \dots x_k, y)] \text{ explains } [x_1 \dots x_k \text{ explain } y]]$

On this proposal, the superficiality of supervenience’s explanatory value comes down to this: the fact that supervenience backs (explains) some explanation is explained by facts involving some deeper explanatory relation. Does this proposal succeed at placing supervenience and paradigmatic explanatory relations on different sides of the Great Explanatory Divide?

I don't think it does. First, let's focus on explanation facts of the form $[[D(w_1 \dots w_n, z)] \text{ explains } [[R(x_1 \dots x_k, y) \text{ explains } [x_1 \dots x_k \text{ explain } y]]]$. In order for Superficial-3 to classify supervenience and grounding in the desired way, the following has to be true: facts of the form $[[y \text{ supervenes on } x_1 \dots x_k] \text{ explains } [x_1 \dots x_k \text{ explain } y]]$ are, whereas facts of the form $[[x_1 \dots x_k \text{ ground } y] \text{ explains } [x_1 \dots x_k \text{ explain } y]]$ aren't, explained by instances of $D(w_1 \dots w_n, z)$. What kinds of explanations are these "meta-explanations"? Presumably, they are non-causal, and more specifically metaphysical, explanations.

This raises the question of what in general explains metaphysical explanations of the form $[[R(x_1 \dots x_k, y) \text{ explains } [x_1 \dots x_k \text{ explain } y]]]$. The best place to look for an answer is the growing literature on iterated ground, whose central question is: "What grounds the grounding facts?".⁴⁸ However, this requires an important bit of qualification. Unfortunately, the literature on iterated ground tends to ignore the distinction between grounding *qua* metaphysical explanation and grounding *qua* the relation underlying metaphysical explanation. In my view, if there are relations underlying metaphysical explanation, the first-order metaphysical explanation facts are explained by facts about these underlying relations. So, the problem of iterated ground is really a problem about what explains *second-order* (and higher up) metaphysical explanation facts: facts about the explanations of the explanations, the explanations of the explanations of the explanations, and so on.

Since Superficial-3 concerns second-order metaphysical explanations of the form $[[R(x_1 \dots x_n, y) \text{ explains } [x_1 \dots x_n \text{ explain } y]]]$, let's focus on these. A glance at the literature reveals four salient options as to what could explain such facts:

⁴⁸ This question was first raised by Sider (2011: 106–111) and Bennett (2011).

(i) *Nothing*. There is no further fact that explains $[[R(x_1 \dots x_k, y)] \text{ explains } [x_1 \dots x_k \text{ explain } y]]$ ⁴⁹

(ii) *The explanans*. $[R(x_1 \dots x_k, y)]$ explains $[[R(x_1 \dots x_k, y)] \text{ explains } [x_1 \dots x_k \text{ explain } y]]$ ⁵⁰

(iii) *Facts about essence*. $[x, z_y]$ explains $[[R(x_1 \dots x_k, y)] \text{ explains } [x_1 \dots x_k \text{ explain } y]]$

(where z_y is some fact about the essence of y)⁵¹

(iv) *Miscellaneous*. It's a miscellany of lower-level facts that explain $[[R(x_1 \dots x_k, y)] \text{ explains } [x_1 \dots x_k \text{ explain } y]]$, but they cannot be subsumed under a general formula⁵²

Keep in mind that for supervenience to be a superficial explanatory relation according to Superficial-3, all facts about supervenience facts explaining explanation facts need to be explained by facts about some deeper explanatory relation distinct from supervenience. But facts about supervenience facts explaining explanation facts are second-order explanation facts, so this necessary condition of superficiality already rules out options (i) and (ii). That is, if either (i) or (ii) is true, then supervenience doesn't satisfy Superficial-3 and is therefore not a superficial explanatory relation.

What about the other two options? (iii) is compatible with supervenience being a superficial explanatory relation; ' $x_1 \dots x_k$ and facts about y 's essence' could mean lots of things, and one of them is essential connectedness between $x_1 \dots x_k$ and y . As I mentioned in

⁴⁹ See Jones ms.

⁵⁰ See Bennett 2011 and deRosset 2013. Litland (2017) argues that the grounding (understood as metaphysical explanation) facts are zero-grounded in the empty set of facts, but emphasizes that his view isn't in competition with the Bennett-deRosset account.

⁵¹ See Dasgupta 2014; cf. Rosen 2010: 130–133 and Audi 2012

⁵² See Sider ms; cf. Wilson 2014 and Author forthcoming-2,

the previous section, some theorists of metaphysical explanation consider essential connectedness a deep explanatory relation, so this would allow us to classify supervenience as only superficially explanatory. The problem now is that by the same reasoning, we would have to classify *all* non-causal explanatory relations as non-explanatory. For according to this interpretation of (iii), all higher-order explanation facts are explained by facts about essential connectedness, and therefore Superficial-3 classifies grounding (*qua* the relation underlying metaphysical explanation) and other paradigmatic explanatory relations as no less superficially explanatory than supervenience. Surely this is not the result advocates of SIN wanted.

Option (iv) allows us to deny that there is any deep explanatory relation uniformly underlying the second-order metaphysical explanation facts and that grounding is only superficially explanatory. Unfortunately for friends of SIN, it allows us to say the same thing about supervenience. If the higher-order metaphysical explanation facts are explained by miscellaneous facts, then we can argue that the explanatory value of supervenience should likewise be explained on a case-by-case basis. Perhaps moral supervenience is explanatory due to conceptual truths about morality⁵³; the supervenience of propositional justification on mental states is explanatory due to facts specific to evidence⁵⁴; etc. The miscellaneous approach leaves many options open, and liberates opponents of SIN from the burden of having to provide a uniform story about the explanatory value of supervenience.

Let's take stock. In this section, I considered a more sophisticated version of the "Surface Relation" worry: perhaps the problem is not that whenever supervenience is explanatory it holds due to other explanatory relations, but that its *explanatory value* is parasitic

⁵³ See Stratton-Lake and Hooker 2006; cf. Zangwill 1997.

⁵⁴ See Conee and Feldman 2004.

on those relations. As I showed, this could mean a number of different things. But on each reading, supervenience either doesn't fall behind the Great Explanatory Divide or also pulls with itself some paradigmatic explanatory relations down into the chasm. This is bad news for advocates of SIN. If the problem with supervenience is that it's only superficially an explanatory relation, but the best sense we can make of this idea is one that leaves it open that grounding, too, might be merely superficially explanatory, then we failed to give the words 'explanatory relation' an interpretation on which SIN comes out both true and interesting.

8. Conclusion

If anything qualifies as a platitude in contemporary metaphysics, the slogan that supervenience is not an explanatory relation certainly does. But insufficient attention has been paid to the question of what the slogan even means. In this paper, I explored a few options but was unable to discern any sense in which the slogan is true but doesn't generalize to some paradigmatic explanatory relations. Now of course, it's difficult to prove a negative, and I cannot rule out some yet-to-be-proposed interpretation of SIN that would make it both true and interesting. But the difficulty of coming up with such an interpretation at least strongly suggests that the slogan doesn't deserve the platitude status it enjoys in contemporary discussions.

Though this is an interesting moral, we shouldn't overstate its significance. I would advise against concluding that supervenience is explanatory. The point is, rather, that as of yet we have no reason to assume a deep, principled difference between supervenience on the one hand, and causation, grounding, realization (etc.) on the other, with respect to their explanatory value; we have been given no useful sense in which supervenience and

paradigmatic explanatory relations fall on different sides of the Great Explanatory Divide. I also said nothing to suggest that anything is wrong with specific objections to supervenience-based formulations of physicalism, moral naturalism, or other (supposedly explanatory) theses. Rather, I think that different explanatory purposes require different relations: supervenience may be unsuitable for capturing the explanatory connection between the mental and the physical, but then again grounding is also unsuitable for capturing the explanatory connection between smoking and lung cancer. And there might be further theses, for example mentalism about propositional justification, for which supervenience is just perfect.

While I had difficulty coming up with a sense in which SIN is both true and interesting, I'm also not sure what in general it would take for a relation to be explanatory. As we have seen above, various criteria have been implicitly assumed in the literature, and it's not clear that any combination of them circumscribes an interesting class of relations that should be regarded as the elite class of explanatory relations. For this reason, I suspect that we would be better off just abandoning the notion of an explanatory relation: there is no sufficiently clear and useful concept in the vicinity that warrants organizing our theories of explanation around it. But this is the topic for another paper.

References

- Audi, Paul (2012), "Grounding: Toward a theory of the *in-virtue-of* relation," *Journal of Philosophy*, 109: 685–711
- Bennett, Karen (2004), "Global Supervenience and Dependence," *Philosophy and Phenomenological Research*, 68: 501–529
- Bennett, Karen (2011), "By Our Bootstraps," *Philosophical Perspectives*, 25: 27–41
- Berker, Selim (forthcoming), "The unity of grounding," *Mind*
- Blackburn, Simon (1985/1993), "Supervenience Revisited," reprinted in *Essays on Quasi-realism*, Oxford: Oxford University Press
- Bliss, Ricki (2014), "Viciousness and Circles of Ground," *Metaphilosophy*, 45: 245–56
- Bliss, Ricki and Trogon, Kelly, "Metaphysical Grounding", *The Stanford Encyclopedia of Philosophy* (Winter 2016 Edition), Edward N. Zalta (ed.), URL = <<http://plato.stanford.edu/archives/win2016/entries/grounding/>>
- Cameron, Ross P. (2014), "Parts Generate the Whole, but They Are Not Identical to It," Donald Baxter and Aaron Cotnoir (eds.) (2014), *Composition as Identity*, Oxford: Oxford University Press, pp. 90–107
- Chalmers, David (1996), *The Conscious Mind*, Oxford: Oxford University Press
- Conee, Earl and Richard Feldman (2004), *Evidentialism: Essays in Epistemology*, Oxford University Press
- Correia, Fabrice (2008), "Ontological Dependence," *Philosophy Compass*, 3: 1013–1032
- Correia, Fabrice (2013), "Metaphysical Grounds and Essence," in M. Hoeltje, B. Schnieder, A. Steinberg (eds.), *Varieties of Dependence*, Philosophia Verlag
- Correia, Fabrice & Benjamin Schnieder (2012), "Grounding: an opinionated introduction," in Correia and Schnieder
- Dasgupta, Shamik (2014), "The possibility of physicalism," *Journal of Philosophy*, 111: 557–592
- Davidson, Donald (1970/1980), "Mental Events," in *Essays on Actions and Events*, Oxford: Clarendon Press
- DePaul, Michael (1987), "Moral supervenience and dependence," *Philosophical Studies*, 51: 425–439
- deRosset, Louis (2010), "Getting Priority Straight," *Philosophical Studies*, 149: 73–97
- deRosset, Louis (2013), "Grounding Explanations," *Philosophers' Imprint*, 13 (7)
- Ehring, Douglas (2011), *Tropes: Properties, Objects, and Mental Causation*, Oxford: Oxford University Press
- Endicott, Ronald (2016), "Functionalism, Superduperfunctionalism, and Physicalism: Lessons from Supervenience," *Synthese*, 193: 2205–2235
- Fine, Kit (2001), "The Question of Realism," *Philosophers' Imprint*, 1 (1)
- Fine, Kit (2012), "Guide to Ground," in Correia and Schnieder

- Grimes, Thomas (1988), "The myth of supervenience," *Pacific Philosophical Quarterly*, 69: 152–60
- Grimes, Thomas (1991), "Supervenience, Dependency, and Determination," *Philosophical Studies*, 62: 81–92
- Guigon, Ghislain (forthcoming), "Truths *qua* Grounds," *Philosophy and Phenomenological Research*
- Hare, R.M. (1952), *The Language of Morals*, Oxford: Clarendon Press
- Hoffman, Vera and Albert Newen (2007), "Supervenience of Extrinsic Properties," *Erkenntnis*, 67: 305–319
- Horgan, Terence (1982), "Supervenience and Microphysics," *Pacific Philosophical Quarterly*, 63: 29–43
- Horgan, Terence (1993), "From Supervenience to Superdupervenience: Meeting the Demands of a Material World," *Mind*, 102: 555–86
- Horgan, Terence and Mark Timmons (1992), "Troubles on Moral Twin Earth," *Synthese*, 92: 221–260
- Jackson, Frank (1998), *From Metaphysics to Ethics*, Oxford: Oxford University Press
- Jackson, Frank and Philip Pettit (1990), "Program Explanation: A General Perspective," *Analysis*, 50: 107–117
- Jones, Nicholas K. (ms), "Ungrounded Ground"
- Kim, Jaegwon (1988/1993), "Supervenience for multiple domains," in *Supervenience and Mind*, Cambridge: Cambridge University Press
- Kim, Jaegwon (1994), "Explanatory Knowledge and Metaphysical Dependence," *Philosophical Issues*, 5: 51–69
- Kim, Jaegwon (1998), *Mind in a Physical World*, Cambridge, MA: MIT Press
- Klagge, James (1988), "Supervenience: Ontological or Ascriptive", *Australasian Journal of Philosophy*, 66: 461–470
- Koslicki, Kathrin (2004), "Constitution and Similarity," *Philosophical Studies*, 117: 327–364
- Koslicki, Kathrin (2013), "Ontological Dependence: An Opinionated Survey," in Miguel Hoeltje, Benjamin Schnieder & Alex Steinberg (eds.), *Varieties of Dependence: Ontological Dependence, Grounding, Supervenience, Response-Dependence*, Philosophia Verlag
- Koslicki, Kathrin (2014), "The Coarse-grainedness of grounding," *Oxford Studies in Metaphysics*, 9: 306–344
- Kovacs, D. M. (2017a). What is wrong with self-grounding? *Erkenntnis*
- Kovacs, D. M. (2017b). Grounding and the argument from explanatoriness. *Philosophical Studies*, 174,
- Kovacs, D. M. (2018). The deflationary theory of ontological dependence. *Philosophical Quarterly*

- Lewis, David (1983), “New Work for a Theory of Universals,” *Australasian Journal of Philosophy*, 61: 343–377
- Lipton, Peter (2001), “What Good is an Explanation?,” in Giora Hon & Sam S. Rakover (eds.), *Explanation: Theoretical Approaches and Application*, Dordrecht: Springer
- Litland, Jon (2017), “Grounding ground”, *Oxford Studies in Metaphysics*, 10: 279–316
- Markosian, Ned (2005), “Against Ontological Fundamentalism,” *Facta Philosophica*, 7: 69–84
- McKenzie, Kerry (2014), “Priority and Particle Physics: Ontic Structural Realism as a Fundamentality Thesis,” *British Journal for the Philosophy of Science*, 65: 353–380.
- McKenzie, Kerry (forthcoming), “Structuralism in the idiom of determination,” *British Journal for the Philosophy of Science*
- McLaughlin, Brian (1995), “Varieties of Supervenience,” in Elias E. Savellos & U. Yalcin (eds.), *Supervenience: New Essays*, Cambridge University Press, 16–59
- McLaughlin, Brian and Bennett, Karen, “Supervenience”, *The Stanford Encyclopedia of Philosophy* (Spring 2014 Edition), Edward N. Zalta (ed.), URL = <<http://plato.stanford.edu/archives/spr2014/entries/supervenience/>>
- Melnyk, Andrew (2003), *A Physicalist Manifesto: Thoroughly Modern Materialism*, Cambridge: Cambridge University Press
- Petrie, Bradford (1987), “Global Supervenience and Reduction,” *Philosophy and Phenomenological Research*, 48: 119–30
- Raven, Michael (2012), “In Defence of Ground”, *Australasian Journal of Philosophy*, 90: 687–701
- Raven, Michael (2015), “Ground,” *Philosophy Compass*, 10: 322–333
- Ridge, Michael (2007), “Anti-reductionism and supervenience,” *Journal of Moral Philosophy*, 4: 330–348
- Rosen, Gideon (2010), “Metaphysical dependence: grounding and reduction,” in Bob Hale and Aviv Hoffman (eds.), *Modality: Metaphysics, Logic, and Epistemology*, Oxford: Oxford University Press
- Ruben, David-Hillel (1990), *Explaining Explanation*, London: Routledge
- Salmon, Wesley (1977/1997), “A Third Dogma of Empiricism,” reprinted in *Causality and Explanation*, Oxford: Oxford University Press
- Schaffer, Jonathan (2009), “On what grounds what,” in Chalmers et al
- Schaffer, Jonathan (2012), “Grounding, Transitivity, and Contrastivity,” in Correia and Schaffer
- Schaffer, Jonathan (2016), “Grounding in the Image of Causation,” *Philosophical Studies*, 173: 49–100
- Schnieder, Benjamin (2010), “A Puzzle About ‘Because’,” *Logique et Analyse*, 53: 317–343
- Schnieder, Benjamin (2015), “The Asymmetry of ‘Because’,” *Grazer Philosophische Studien*, 91: 131–164

- Schnieder, Benjamin (forthcoming), “Grounding and Dependence,” *Synthese*
- Shagrir, Oron (2002), “Global Supervenience, Coincident Entities and Anti-individualism,” *Philosophical Studies*, 109: 171–196
- Shagrir, Oron and Vera Hoffman-Kolss (2013), “Supervenience,” in: B. Kaldis (ed.), *Encyclopaedia of Philosophy and the Social Sciences*, SAGE Publications
- Shoemaker, Sydney (2007), *Physical Realization*, Oxford: Oxford University Press
- Sider, Theodore (1999), “Global Supervenience and Identity Across Times and Worlds,” *Philosophy and Phenomenological Research*, 59: 913–37
- Sider, Theodore (2011), *Writing the Book of the World*, Oxford: Clarendon Press
- Sider, Theodore (ms), “Ground grounded”
- Skiles, Alexander (2015), “Against Grounding Necessitarianism,” *Erkenntnis*, 80: 717–51
- Skow, Bradford (2016a), “Against Understanding (as a Condition on Explanation),” in Stephen Grimm (ed.), *Making Sense of the World: New Essays on the Philosophy of Understanding*, Oxford University Press
- Skow, Bradford (2016b), *Reasons Why*, Oxford: Oxford University Press
- Sosa, Ernest (1991), *Knowledge in Perspective*, Cambridge, Cambridge University Press
- Stratton-Lake, Phillip and Brad Hooker (2006), “Scanlon vs. Moore on Goodness,” in Terence Horgan and Mark Timmons (eds.), *Metaethics after Moore*, Oxford: Oxford University Press
- Strawson, Peter F. (1985), “Causation and Explanation,” in B. Vermazen and M. Hintikka (eds), *Essays on Davidson: Actions and Events*, Oxford University Press
- Strevens, Michael (2008), *Depth: An Account of Scientific Explanation*, Cambridge, MA: Harvard University Press
- Tahko, Tuomas E. and Lowe, E. Jonathan, “Ontological Dependence”, *The Stanford Encyclopedia of Philosophy* (Spring 2015 Edition), Edward N. Zalta (ed.), URL = <http://plato.stanford.edu/archives/spr2015/entries/dependence-ontological/>
- Trogon, Kelly (2013), “An Introduction to Grounding,” in Hoeltje et al
- Trogon, Kelly (forthcoming), “Grounding-Mechanical Explanation,” *Philosophical Studies*
- Van Cleve (1985), “Epistemic Supervenience and the Circle of Belief,” *Monist*, 68: 90–104
- Wilson, Jessica (1999), “How super-duper does a physicalist supervenience need to be,” *Philosophical Quarterly*, 50: 33–52
- Wilson, Jessica (2014), “No Work for a Theory of Grounding,” *Inquiry*, 57: 1–45
- Yablo, Stephen (1992), “Mental Causation,” *Philosophical Review*, 101: 245–280
- Zangwill, Nick, (1997), “Explaining Supervenience: Moral and Mental,” *Journal of Philosophical Research*, 22: 509–18