Metaphysical Emergence as Higher-Level Naturalness

Alastair Wilson
University of Leeds & Monash University
a.j.j.wilson@leeds.ac.uk

14th March 2024

Abstract: In this chapter I explore an approach to metaphysical emergence which distinguishes between fundamentality and naturalness and endorsing the thesis that there are natural properties at non-fundamental levels. I take as my starting point Elizabeth Barnes's proposal to characterize the emergent as fundamental but dependent, criticizing it on the ground that it undermines the theoretical work we need fundamentality to do. However, I think Barnes is on the right track: emergence is linked to a selective metaphysical privileging of higher-level subject-matters. I suggest an alternative account of the metaphysically emergent as non-fundamental but (at least relatively) *natural*, and show how this suggestion can be implemented in a simple subject-matter-based framework.

1. Introduction

In this chapter I outline a new approach to the metaphysics of emergence which works by distinguishing between fundamentality and naturalness and endorsing the thesis that there are natural properties at non-fundamental levels. I take as my starting point Elizabeth Barnes' proposal (Barnes 2012) to characterize the emergent as dependent but fundamental, criticizing it on the grounds that it undermines the theoretical work we need fundamentality to do. However, I think Barnes is on the right track: emergence is linked to a selective metaphysical privileging of higher-level subject-matters, although fundamentality is not the best way to think about that privileging. I suggest an alternative account of the emergent as non-fundamental but (at least relatively) *natural.* In a slogan: metaphysical emergence is higher-level eliteness. I conclude by sketching a formalization of the account, drawing on Russell (2015).

2. Factorizing Emergence

Assume for simplicity a layer of non-emergent fundamental facts as a background to our discussion. These may be facts about particles, fields, spacetime points, or the like. Call these facts the fundamental base. Against this background, what is it to emerge? A natural thought is that something emerges when a higher-level phenomenon depends on an underlying-level phenomenon, but in a way which introduces some distinctive or novel feature. So I will start by factorizing emergence into two components: dependence and distinctness. The emergent is that which is both suitably dependent on and suitably distinct from facts in the fundamental base.

The dependence component of our factorization is that the emergent is dependent on the fundamental base, at least in the minimal sense of global supervenience. It is not metaphysically possible for the emergent phenomena to vary without some variation or other in the fundamental base. Some theories of strong emergence involve denying supervenience, but I won't question it here: my target is a metaphysically robust notion of emergence that still respects supervenience.

The distinctness component of our factorization is that the emergent is distinct in some significant way from the fundamental base. Often it is said for example that the emergent is *something over and above* the basis. But what sort of distinctness is involved? The factorization permits a natural distinction between different kinds of emergence, based on different ways in which the higher-level phenomenon might be distinct from the underlying level. For example, some epistemic or explanatory accounts of distinctness give rise to 'weak' or 'explanatory' forms of emergence. However, our focus here is on more metaphysical accounts of distinctness, and the corresponding metaphysically robust forms of emergence that result.

We can now contrast various different ways in which a higher-level phenomenon might be metaphysically distinct from the lower-level phenomena on which it depends. Preliminary options include: the emergent is distinct by being *novel* when its dependence base is not, the emergent is distinct by being *autonomous* when its dependence base is not, the emergent is distinct by being *downwards causal* when its dependence base is not, or the emergent is distinct by being fundamental when its dependence base is not.

I will suggest an alternative view: the emergent is distinct by being *higher-level (perfectly) natural* when its dependence base is not.

3. Emergence vs Fundamentality

In order to develop this chapter's positive account, I will draw on the influential recent discussion of Barnes (2012) as a foil. Barnes argues that emergent phenomena are well characterized as fundamental despite being dependent:

> "emergent entities look plausibly characterized as fundamental. They are part of the real stuff, the core stuff, in the world. They are not the stuff you get for free." - Barnes (2012), p.885

Barnes elaborates this line of thought with two informal 'glosses' on fundamentality, which are intended to bring out the sense in which the emergent adds something fundamental to the fundamental dependence base. The first gloss is that when God creates the world, they have to create the emergent as well as the base. The second gloss is that the base does not contain truthmakers for all truths about the emergent.

We should bear in mind that Barnes's two glosses on the theoretical role of fundamentality are intended to communicate a way of thinking, rather than to persuade a dug-in opponent. But even acknowledging that, it is hard to see how the glosses can be reconciled with a (the?) core role of fundamentality: modally fixing – necessitating – the non-fundamental. Consider the two glosses in turn. Given that the emergent emerges from the base, then a defender of the traditional view is in a position to maintain that – after all – all that God has to do is to make the base. The emergent will – by metaphysical necessity! – emerge from it. Likewise, if truths about the emergent globally supervene on the base, then a defender of the traditional view can point out that the base already contains enough to modally fix the emergent. Once that modal fixing is acknowledged, what sense could it make to deny that the base truth-makes all truths about the emergent?

In light of the response that emphasizes modal fixity, it is not difficult for defenders of an orthodox account of fundamentality to consistently resist Barnes's line of thought. They can start by maintaining that non-fundamental entities are real. (Only compositional nihilists deny their existence, and that view is implausible.) Next, are non-fundamental entities core stuff? Well, if 'core stuff' means 'fundamental stuff', the argument straightforwardly begs the question. If, on the other hand, 'core stuff' means 'important stuff', then, since the derivative is important, it can be core stuff. Finally, do we get emergent stuff for free? Well, on the orthodox approach, we do pay for the emergent, but wholly in the coin of the base.

What does this line of response on behalf of the orthodox account show us? Dialectically, what matters is not that some response or other be available (there is always a premise which can be denied in any argument!) but rather that the response be in the spirit of the orthodox account – that it be a principled reply rather than an ad hoc fix. And the response does seem principled enough to me. So far, then, I don't think Barnes has identified any tension in the orthodox view. Distinguishing independence from fundamentality seems to be a dead end.

A further clue to a potential way forward, however, can be found in the truthmaker gloss. Barnes argues that an absence of truthmaking by the base, even in the presence of modal fixing, suffices for the emergent to count as fundamental. But why should this be so? I hypothesise that it is related to a key element in truthmaking, which takes it beyond mere modal fixing of truth: *aboutness*. A variant of Barnes' line of thought might emphasize this point: the emergent is fundamental because no facts in the base are *about* the emergent, even though the base modally fixes the emergent. I would firmly deny that this lack of aboutness justifies our treating the emergent as fundamental in the manner which Barnes advocates. But I do think that the lack of aboutness may point us towards a different account of what is distinctive about the emergent.

To characterise the metaphysical distinctiveness of the emergent relative to its dependence base, we can look beyond fundamentality and substitute a different concept in the vicinity: *naturalness*. Naturalness can be thought of as a higher-order property, a property of properties, which confers theoretical significance on its bearers. How exactly that significance is cashed out varies amongst different proponents of naturalness. Natural properties are variously said to be elite, privileged, or non-gerrymandered, to carve nature at the joints, or to draw distinctions corresponding to genuine differences. Detailed theories of naturalness differ on various specific points, tailored to differences in their intended applications; these differences will be discussed in section 4. My intended notion of naturalness broadly follows Quine (1969), Lewis (1983), Schaffer (2004), Dorr & Hawthorne (2013), and Thompson (2016). The application of naturalness I will be making is reasonably tolerant of differences between conceptions of naturalness. However, one distinctive aspect of naturalness is non-negotiable if things are to get off the ground. In the Lewisian system, (perfect) naturalness also plays the role of fundamentality; here I am suggesting pulling these two notions apart. This manoeuvre makes conceptual room for the non-fundamental perfectly natural properties which will be our candidates for emergent properties.[1]

---

[1] It also in principle opens the way for fundamental properties which are not perfectly natural. Even if a detailed theory permits such properties, though, they seem unlikely to have much explanatory significance.

4. The Fundamentality Role and the Naturalness Role

In this section, I will fill out the details of my suggestion to distinguish fundamentality from perfect naturalness. First, I will show how the different aspects of the naturalness role as characterised by Lewis, Sider and others can be disentangled, so that some of them are played by a concept of fundamentality and some of them are played by a concept of naturalness. Next I look at some previous treatments of naturalness at non-fundamental levels, including by Schaffer and List, and I show how these treatments can be co-opted in service of the present account of emergence. This paves the way for a schematic formalization of the account in section 5.

Although it is contemporary orthodoxy within the metaphysics of properties that the fundamental properties are perfectly natural, and vice versa, I think it makes good theoretical sense for these statuses to come apart. Intuitively, fundamentality and naturalness do different jobs in our metaphysical ideology. Fundamentality (absolute and relative) tells us what depends on what: the less dependent, the more fundamental. Naturalness (perfect and imperfect) tells us the optimal way to carve up the whatnots: the better the carving, the more natural. So understood, absolute fundamentality and perfect naturalness only coincide if the most fundamental description of the world is the unique best carving; but that coincidence is exactly what is denied by the kind of metaphysical emergentist I have in mind. (Perhaps it is metaphysically possible for all perfectly natural properties to reside at the fundamental level; that would amount to a degenerate case of the framework presented in the next section. But our immediate focus is on what is actual.)

The fundamentality role has at its core the modal connection discussed above: the fundamental facts are supposed to form a minimal supervenience base upon which all the other facts are obliged to supervene. The fundamental facts, on this view, are the smallest subplurality of facts such that they collectively fix all the facts. Relatedly, fundamentality is employed to account for duplication of particular individuals: two individuals are duplicates if and only if they match with respect to all of their fundamental properties.

Fundamentality comes in both absolute and relative forms. There is debate about whether absolute fundamentality can be reduced to relative fundamentality, or about whether there even is anything which is absolutely fundamental; but, in any case, we need some relative notion or other if we are to make serious use of the notion in interlevel metaphysics. Relative fundamentality generates a partial ordering; this point extends the connection with supervenience just discussed, since (one-way) supervenience also generates a partial ordering. In general, a less fundamental description of some phenomenon supervenes on a more fundamental description.

The naturalness role is more disputed than the fundamentality role. Hawthorne and Dorr (2013) identify numerous components which have at one stage or another been attributed to the naturalness role. For present purposes, I think the following three points are crucial.

The first point is that natural properties have a special connection to reference via eligibility: a theory which includes predicates and other terms that refer to natural properties is thereby a better theory. In other words, we ought to theorise in natural terms. A broadly interpretational metasemantics transmutes this betterness into the phenomenon of 'reference magnetism': our terms are attracted, as if by a magnet, to refer to the most natural candidate referents in the vicinity. This phenomenon is associated with a popular solution to 'Putnam's paradox' (Putnam 1977, Lewis 1983) and related permutation-based sceptical arguments against forms of metaphysical realism. The general idea is that permutations of plausible reference assignments which preserve truth but which associate linguistic predicates with highly non-natural properties can be disregarded as candidate interpretations.

Second (and closely related to the first point), natural properties are supposed to be especially eligible to serve as the contents of our concepts. A thought which involves concepts corresponding to more natural concepts is all other things being equal a better thought. In other words, we ought to think in natural terms. A broadly interpretational view of mental content then enables us to bring natural properties to bear to solve Kripke's version of Wittgenstein's rule-following paradox (Kripke 1982): roughly, we can rightly interpret a reasoner as using addition rather than the deviant variant quaddition because addition, but not quaddition, picks out a relatively natural mathematical function.

Finally, natural properties are supposed to be especially projectible, in the sense of being especially eligible to feature in successful inductive inference. This seems to me the least secure of the three roles, at least in a wider epistemological and metaphysical context – for example, it is just not clear how natural properties count as more projectible in the Lewisian system, given the contingency of laws governing perfectly natural properties in that system. However, my main purpose here is not to refine accounts of naturalness but to explore their application.

Like fundamentality, naturalness is typically supposed to come in absolute and relative versions. Even if 'purple' is not perfectly natural, 'purple' is still more natural than 'purple or ingratiating'. will also be a feature of the proposal of the present chapter: naturalness, as linked to theory-worthiness, is the kind of thing which can be had to a greater or lesser extent. This feature of naturalness, allied with this chapter's account of the emergent as natural but non-fundamental,

has the result that emergence itself is a matter of degree. For example, if both a long-lived standing wave pattern and a short-lived pattern which momentarily looks like a fish both appear in a dish of vibrating sand, then the latter pattern – because less natural – counts as *less emergent* then the former.

Granting the above characterizations – of fundamentality as a matter of dependence and independence and of naturalness as a matter of optimal theoretical description – then Barnes's approach looks misplaced. Naturalness appears eminently better suited than fundamentality to characterizing the distinctness of emergent properties relative to their dependence base: emergent phenomena do seem to be optimally carved up in terms of properties at their own level, rather than in terms of properties at any more fundamental level. And (as I argued in section 3) the way the emergent is modally fixed by its base tells against regarding it as fundamental. Still, however, I think the present proposal has something important in common with Barnes's account: emergent properties are genuinely metaphysically privileged relative to their dependence bases, by being more natural than them.

For this chapter's account of emergence as higher-level naturalness to get any traction, there must be natural properties at multiple levels of reality – not just at the fundamental level. This marks a deviation from Lewis, who acknowledges perfect naturalness at the fundamental level only, and from Sider, who follows him in this. Although both authors allow for a notion of relative naturalness, they conceive it as something which strictly decreases up the modal hierarchy of levels: the supervenient can never be more natural than its supervenience base. This feature is ultimately a consequence of the conflation of naturalness with fundamentality in the Lewisian tradition.

By contrast, however, Schaffer (2004) defends a conception of naturalness attaching to higher-level properties which he calls the 'scientific conception'. Schaffer's view is well-suited to characterizing emergence as high-level naturalness. Emergent high-level phenomena are bearers of natural properties, while non-emergent high-level phenomena do not bear natural properties.

An immediate concern about the scientific conception of natural properties is whether it really obtains any support from the discoveries or the practice of science. What is it about natural properties which make them scientifically relevant? It's not as if scientists have successfully constructed a naturalness detector! In that case, what does it matter whether some higher-level phenomenon is natural (and hence metaphysically emergent) or unnatural (and hence not)?

The answer to this challenge, however, is already implicit in existing naturalness proposals. As already discussed, naturalness connects with scientific theorizing in two related ways: through projectibility of scientific predicates (addressing Goodman's New Riddle of Induction), and through reference-magnetism for our scientific concepts (addressing Putnam's Paradox and Kripkenstein's puzzle.) There are, of course, further questions to be asked about theoretical posits in fundamental metaphysics function; some might say that we need a positive argument that naturalness can play the naturalness role. Here, though, I would tend to follow Schaffer (2016) in rejecting the demand at least for situations like the present one.

Consider some examples to get a feel for the present proposal. First I will suggest some examples involving candidate cases of perfect naturalness. Phenomenal properties might be perfectly natural; then conscious mental phenomena could be metaphysically emergent even if they are wholly dependent on a microphysical fundamental base. Geometric structural properties might be perfectly natural; then chemical phenomena involving molecular structure could be metaphysically emergent even if they are wholly dependent on a quantum-field-theoretic fundamental base. Spatiotemporal relations might be perfectly natural; then spacetime could be metaphysically emergent even if it is wholly dependent on a non-spatiotemporal fundamental base.

Next, consider some examples involving candidate cases of imperfect naturalness. Colour properties might be imperfectly natural – since some contribution from contingent human physiology and culture is involved in defining them – but colour properties can still be emergent so long as they are more natural than their microphysical dependence base. Thermodynamic collective quantities like temperature and pressure might be imperfectly natural – for example insofar as there is some arbitrariness in their definitions or extension to borderline cases like black holes – but they can still be emergent if they are more natural than the microphysical state descriptions in terms of particle positions and momenta on which they supervene. And folk attributions of mental states like anger and envy might be imperfectly natural – for example if there is psychological or cultural contingency in our folk mental categories – yet remain more natural than the patterns of neuronic activity which subvene them.

In this section I aim to have presented the initial features of the account of emergence as higher-level naturalness, and in particular to have exhibited how it relies on an underlying account of naturalness which does not monotonically decrease as we move up the levels of nature. In the next section, I set out one such account of naturalness which is hospitable to emergence as higher-level naturalness.

5. A Subject-Matter Account of High-Level Naturalness

It would be nice to have a detailed account of higher-level naturalness to underwrite my schematic account of emergence. To fill this gap, I will exhibit a basic construction of naturalness in terms of the notion of a subject-matter, drawing on Lewis (1988) and Russell (2015). Fundamentality and naturalness can be unified in that they both generate metaphysically interesting orderings of subject-matters; these orderings can also interact in interesting interesting ways.

In the framework I will propose, fundamentality is understood in essentially modal terms. A more fundamental subject matter is just one which draws more fine-grained distinctions amongst possible worlds, so that the most fundamental description is also the maximally detailed description. Perfect naturalness, by contrast, is a primitive predicate which attaches to partitions of possible worlds. Perfect naturalness attaches to the finest-grained partition (the 'fundamental level'), but also to certain coarse-grained partitions ('higher levels'). Subject-matters can also possess imperfect naturalness to varying degrees, ranging from highly but not perfectly natural properties (consider: the defined but theoretically useful terms within some scientific theory) to properties with a much lower degree of naturalness (consider: gerrymandered miscellaneous properties like being a comma or a cosmonaut).

Here are some definitions of key terms in the construction to follow:

Subject-matters: A subject-matter is a partition on a set of worlds W, i.e. a set of subsets of W such that their union is W and the intersection of any pair of them is empty.

Supervenience of subject-matters: A supervenes on B iff all worlds which are in the same cell of B are in the same cell of A.

Inclusion of subject-matters: C includes D iff D supervenes on C.

Proper inclusion of subject-matters: E properly includes F iff E includes F and F does not include E.

Aboutness: Proposition P is about a subject-matter S if P supervenes on S.

Exact aboutness: Proposition P is exactly about a subject-matter S iff P is about S, any other subject-matter that P is about is included in S, and there is no other subject-matter of which the same is true.

With those definitions in place, the construction proceeds as follows:

1. Start with any language L, which is capable of describing reality completely (though potentially with redundancy).
2. Construct a set of worlds W_L as maximal consistent sets of sentences (ultrafilters) in L.
3. Construct the subject-matter U_L which places every world in W_L into a cell of its own.
4. Let U_L include a metaphysically privileged subject-matter N: how the fundamental level of reality naturally is.
5. Construct a set of natural worlds W_N as equivalence classes of members of W_L under the relation: equivalence with respect to N.
6. Construct a language L_N such that each natural world in W_N is a maximal consistent set of sentences in L_N, and vice versa.

L_N is exactly expressively adequate to describe the fundamental subject-matter. If God wrote the book of the world (cf. Sider 2011), he'd write it in L_N.

7. Define relative fundamentality of factual subject-matters: A is more fundamental than B iff A properly includes B.
8. Let there be privileged *perfectly natural* subject-matters N_i such that N is more fundamental than each of the N_i.

With the notion of a perfectly natural subject matter in hand, we can identify an emergent level of nature as one of the N_i, and an emergent proposition as one which is exactly about an emergent subject matter.

The higher-level natural subject matters N_i form a partial order. Though we do obtain a hierarchical picture of reality where fundamentality strictly decreases as one moves upwards, the resulting picture is more like a web of levels than a ladder of levels. The theory of ontological levels recently developed by List (2018) can be adapted to give a formal account of this partial ordering; List's 'ontological' plays essentially the same theoretical role as my 'perfectly natural'.

The next step is to move from the notion of an emergent subject matter to the more familiar level of emergent entities and properties. Here, the flexibility of the subject-matter approach is manifest in that there are various ways to think about the decomposition of some emergent subject-matter into emergent entities and emergent properties depending on what sort of metaphysics of properties is preferred. Regardless of what specific account we prefer, though, we

can identify emergent properties as whatever properties are involved in our preferred object-property decomposition of an emergent level, and emergent entities as whichever entities are the bearers of those emergent properties. Using the notion of subject-matter in this way allows us to sidestep some difficult questions about the metaphysics of properties. As an example: posing the account in terms of natural subject-matters allows us to avoid arbitrary choices among symmetric sets of properties, without introducing redundancy. Consider Kit Fine's example (Fine 2001) of three mutually dependent properties: mass, density, volume in a homogenous fluid. In this scenario, any two properties suffice to determine the third. Taking individual properties as the primary bearers of fundamentality and naturalness thereby leads to potential arbitrariness: we can identify any two of the three as fundamental and perfectly natural, with the other one derivative, but it doesn't matter which two we choose.

Defining fundamentality for properties is, by contrast, relatively straightforward. We obtain an account of fundamentality for properties as follows: first define the factual subject-matter $S(P)$ of a property P as the partition such that all worlds in each subset agree on the distribution of P. (Since we are trying to stay neutral between specific accounts of properties, we will treat the notion of a property-distribution as basic for present purposes.) Next define the level $L(P)$ of a property P as the least fundamental N_i which includes $S(P)$. Then define relative fundamentality of properties: P is more fundamental than Q if $L(P)$ properly includes $L(Q)$. Since the relation of proper inclusion is irreflexive, transitive and anti-symmetric, relative fundamentality so conceived is a partial ordering of subject-matters, and it generates a partial ordering of properties.

The present account locates a property P at the least fundamental level which needs to be specified in order to specify the distribution of P. For arbitrary truth-functional combinations of properties $F(P_i)$, the account says that $L(F(P_i))$ is the least fundamental N_i that includes all the $L(P_i)$. This may seem strange: electron or hippopotamus comes out as more fundamental than molecule. But it is well-motivated: we need to specify reality in more detail to fix the distribution of electrons than we need to in order to fix the distribution of molecules; a fortiori, we need to specify reality in more detail to fix the distribution of electrons and hippopotami than we need to in order to fix the distribution of molecules. This point brings out the distinction between the fundamentality of a property and its naturalness. Being a tree is not a fundamental property, but its distribution is specifiable at the botanical level in a straightforward way. In contrast, being a tree or a tritium atom is a much more fundamental property, since its distribution can only be specified at the microphysical level; however, that microphysical specification is overwhelmingly complicated and in practice unavailable. When it comes to usability, naturalness is what matters, not fundamentality.

What has all this got to do with grounding, and the account of fundamentality as ungroundedness which is its hallmark? Well, the notion of fundamentality associated with grounding is clearly conceptually distinct from the modal notion operative in this chapter. In orthodox grounding theory (Schaffer 2009, Bennett 2017) fundamentality is understood as ungroundedness (and indeed this account of fundamentality is often regarded as an important theoretical application for the notion of grounding). In the present approach, fundamentality is understood as modal fixing of everything. However, these conceptions could potentially be brought into alignment, across all possible cases where they overlap, through the addition of an auxiliary assumption: the key linking principle would be that *whenever A (non-trivially) modally fixes B, being such that A is a way of being such that B*. If this principle holds, then non-trivial modal fixing coincides with grounding amongst contingent state of affairs. Modal fundamentality can be seen as a restricted notion of fundamentality: fundamentality in the contingent domain. The proposal of this chapter then becomes: emergence is naturalness plus *modal* non-fundamentality.

It is worth noting that the proposed approach to emergence is not in principle tied to synchronic emergence, so that the door remains open to the kind of diachronic emergence that is defended by Humphreys (this volume) and criticised by J. Wilson (this volume). Given that the fundamental is characterised in this chapter in terms of modal fixing, diachronic emergence can potentially be understood in terms of (more fundamental) facts at one time metaphysically necessitating (natural although less fundamental) facts at another. That possibility might of course be inconsistent with principles of recombination of various kinds, but we need not take a stand on those here. For more discussion of the relation between kinds of dependence and the diachronic/synchronic/achronic distinction, see A. Wilson (2020).

The present proposal intentionally leaves open most questions about the relation between causation and emergence. Since it is plausible that causation can in some cases link higher and lower-level facts – for example when a single nuclear decay triggers an explosive chain reaction in fissile material – then emergent properties, conceived of as higher-level natural, can in principle cause and be caused. Indeed, on most accounts of naturalness, facts about naturalness are themselves partially responsible for the facts about what caused what: it is plausibly at least in part because a chicken is more natural than a chicken's front half that we say we were pecked by a chicken rather than by the chicken's front half. This renders the account in principle compatible with accounts which link emergence with downward causal role, such as the contributions to this volume by Drossel, Gillett, Hendry, Mitchell, Silberstein, and Yates.

Could a stronger case might be made here? – does the higher-level naturalness view make downwards causation likely, or even inevitable? Let us assume that the higher levels will involve natural higher-level individuals standing in some pattern of natural higher-level properties (although no such breakdown into individuals and properties has been set out in the chapter so far). Let us further assume that these natural higher-level individuals will in general stand in higher-level causal relations. It remains open to a defender of higher-level naturalness to continue to insist that causal relations never point downwards; that is to say, nothing ever causes anything at any level below itself. Perhaps causal relations might be assumed to be always horizontal, or perhaps causal relations might be allowed to go up the levels as in the chain fission reaction example. Although these pictures would each need further justification, for present purposes what I'd like to observe is that there is no obvious incoherence in them. Higher-level naturalness, while it is hospitable to the existence of higher-level causes and effects, does not incorporate any restrictions on the directionality of the causal relations which hold between them.

I conclude that while the higher-level naturalness view makes room for downwards causation, it does not obligate any particular pattern for the causal relations. That is, I think, a desirable result in what is – so far, at least – a theory framework rather than a detailed theory of emergence. A general framework for emergence should be, where possible, congenial to a wide variety of hypothesized explanatory patterns that might be explored as candidate accounts of some puzzling phenomenon. This is not to say that we must regard all patterns of causation across higher-level natural subject matters as metaphysically possible; given a broadly abductive approach to metaphysical modality, it is plausible that if causation is in fact non-downward then causation is non-downward as a matter of metaphysical necessity. But we should, at least initially, regard all patterns of causal relations amongst higher-level subject-matters as epistemically possible for the purposes of doing our metaphysics.

6. Conclusion

I have argued that what makes an emergent phenomenon distinct from its fundamental basis is not fundamentality but naturalness. Naturalness, if distinguished from fundamentality and permitted to attach to higher levels, can be used to characterize emergence. Naturalness orders descriptions of the world with respect to how well they carve reality at its joints. Fundamentality orders descriptions of the world with respect to how much else they modally fix. Then to be emergent is to be non-fundamental but natural.

My schematic account leaves many questions open. Are the facts about high-level naturalness themselves grounded in fundamental facts? Is this grounding connection opaque, and if so does that undermine the thought that the high-level ones are genuinely natural? More generally: are facts about perfect naturalness fundamental? Are facts about fundamentality perfectly natural? Analogues of Sider's Purity principle – that fundamental truths contain only fundamental notions – can be formulated to settle these questions; however, it is outside the scope of the present chapter to evaluate them. For now, the main conclusion of the chapter is that the account of emergence as higher-level naturalness deserves further attention from metaphysicians. It does justice to the intuitive motivations for emergence, and uses theoretical tools which have independent application in the theory of properties. For theorists who already recognize a distinction between fundamentality and naturalness, so that naturalness can accrue to higher-level properties, all the ingredients needed to make the account work are already in place. For theorists who recognise only one notion in the vicinity of fundamentality and naturalness, some ideological overhaul will be required. This overhaul could be partially justified in terms of enabling an interesting theory of emergence, although I argued in section 4 that it has some independent plausibility.

In any case, I hope to have added to the list of theoretical options for understanding the metaphysics of metaphysical emergence.[2]

References

Barnes, E. (2012). "Emergence and Fundamentality", *Mind* 121(484): 873-901.

Bennett, K. (2017). *Making Things Up.* Oxford: Oxford University Press.

Dorr, C. & Hawthorne, J. (2013). "Naturalness", *Oxford Studies in Metaphysics* 1. Oxford: Oxford University Press.

Fine, K. (2001). "The question of realism", *Philosophers' Imprint* 1(2): 1-30.

Kripke, S. (1982). *Wittgenstein on Rules and Private Language: An Elementary Exposition.* Cambridge: Harvard University Press.

Langton, R & Lewis, D. (1998). "Defining "intrinsic"", *Philosophy & Phenomenological Research*, 58(2): 333-345.

Lewis, D. (1983). "New work for a theory of universals", *Australasian Journal of Philosophy* 61: 343-377.

Lewis, D. (1988), "Sentences Partly About Observation", *Philosophical Papers* 17(1): 1-31.

List, C. (2019). "Levels: Descriptive, Explanatory, and Ontological", *Noûs* 53(4): 852-883.

Putnam, H. (1977). "Realism and Reason", *Proceedings of the American Philosophical Association* 50: 483-498.

Quine, W. V. O. (1969). "Natural Kinds", in Jaegwon Kim & Ernest Sosa (eds.), *Ontological Relativity and Other Essays.* Columbia University Press. pp. 114-38.

Russell, J. (2015). "Possible Worlds and the Objective World", *Philosophy & Phenemenological Research* 90(2): 389-422.

Schaffer, J. (2009). 'On What Grounds What', in *Metametaphysics: New Essays on the Foundations of Ontology,* eds. D. Chalmers, D. Manley & R. Wasserman. New York: Oxford University Press.

Schaffer, J. (2004). "Two conceptions of sparse properties", *Pacific Philosophical Quarterly* 85: 92-102.

Schaffer, J. (2016). "It is the Business of Laws to Govern", *dialectica* 70(4): 577-588.

Sider, T. (2011). *Writing the Book of the World.* Oxford: Oxford University Press.

Thompson, N. (2016). "Is Naturalness Natural?", *American Philosophical Quarterly* 53(4): 381-395.

Wilson, A. (2020). "Classifying Dependencies", in *The Foundation of Reality: Fundamentality, Space and Time*, eds. D. Glick, G. Darby and A. Marmodoro. Oxford University Press, 2020.