Metaphysically explanatory unification

David Mark Kovacs[1]

Abstract This paper develops and motivates a unification theory of metaphysical explanation, or as I will call it, Metaphysical Unificationism. The theory's main inspiration is the unification account of scientific explanation, according to which explanatoriness is a holistic feature of theories that derive a large number of explananda from a meager set of explanantia, using a small number of argument patterns. In developing Metaphysical Unificationism, I will point out that it has a number of interesting (and to my mind, attractive) consequences. The view offers a novel conception of metaphysical explanation that doesn't rely on the notion of a "determinative" or "explanatory" relation; it allows us to draw a principled distinction between metaphysical and scientific explanations; it implies that naturalness and fundamentality are distinct but intimately related notions; and perhaps most importantly, it re-establishes the unduly neglected link between explanation and understanding in the metaphysical realm. A number of objections can be raised against the view, but I will argue that none of these is conclusive. The upshot is that Metaphysical Unificationism provides a powerful and hitherto overlooked alternative to extant theories of metaphysical explanation.

Keywords Grounding · Metaphysical explanation · Naturalness · Scientific explanation · Unification theory of explanation

✉ David Mark Kovacs
  david.mark.kovacs@gmail.com

[1] Department of Philosophy, Tel Aviv University, Gilman 363, 6997801 Ramat Aviv, Tel Aviv, Israel

1 Introduction

Philosophers are constantly in the business of offering explanations. Material objects, they say, owe their features to those of their parts; the singleton set {Socrates} exists because Socrates does; torture is wrong because it doesn't maximize utility; we feel pain because our C-fibers are firing. Explanations like these, which philosophers frequently call *metaphysical* explanations to distinguish them from causal explanations, are everywhere. Since 'metaphysical explanation' is a philosophers' term of art, I can't neutrally characterize it. Instead, I will defend my own *unification account*, according to which metaphysical explanation is a matter of systematizing the phenomena in a certain way.

Most metaphysicians think of metaphysical explanation against the backdrop of what I will call the *Backing Model*. Proponents of this model think that explanations work because of "explanatory" or "determinative" relations that "back" them.[1] The relevant relation is typically thought to be causation in scientific explanations and something else in metaphysical explanation. What is this "something else"? "Separatists" think it's *grounding*, a metaphysical analogue of causation, whereas "unionists" reserve the word 'grounding' for metaphysical explanation itself [I borrow this terminology from Raven (2015)]. And "strong grounding pluralists", as Berker (2018) calls them, eschew 'grounding'-talk altogether and posit a multitude of relations underlying metaphysical explanation.[2] Note that while unionists and strong grounding pluralists reject the idea that grounding backs metaphysical explanation, they don't thereby claim that nothing backs it. They typically appeal to other relations, usually essential connectedness in the case of unionists[3] and composition, realization, micro-based determination (etc.) in the case of strong grounding pluralists.[4] By contrast, the view I will develop this paper dispenses with explanatory relations altogether.[5]

While most theories of metaphysical explanation presuppose the Backing Model, non-backing views have been prominent in the scientific explanation literature. Pragmatic theories identify explanations with speech act types or logical constructions thereof.[6] The largely abandoned covering-law model takes explanations to be

---

[1] See Ruben (1990), Kim (1994), Audi (2012) and Schaffer (2016) for and Kovacs (2017) and Taylor (2018) against the Backing Model. The view is also widely referred to as "explanatory realism", I think misleadingly: as I will later show, one can be a realist about explanation without endorsing the Backing Model.

[2] Unionists include Fine (2001, 2012), Dasgupta (2014a, b), Raven (2015) and Litland (2017). Most extant critiques of grounding are best interpreted as only targeting grounding as a "backing" relation, not grounding simply understood as metaphysical explanation (Daly 2012: 89, 94; Wilson 2014: 539 l; Koslicki 2015).

[3] Rosen (2010), Correia (2013), Fine (2015) and Litland (2017).

[4] Wilson (2014) and cf. Bennett (2017).

[5] By this, I don't mean to rule out "explanatory relations" downstream from the notion of metaphysical explanation. DeRosset (2019: Ch. 1) argues that given a suitably abundant conception of facts, grounding as a relation between facts is cheap (cf. Kovacs 2017: 2936, 2937). See also Thompson (2016) and Dasgupta (2017) for attempts to get rid of backing relations.

[6] van Fraassen (1980) and Achinstein (1983).

arguments with no redundant premises and with at least one premise that states a law of nature.[7] And the unification view (or briefly, *Scientific Unificationism*) treats explanatoriness as a holistic feature of theories that derive as many explananda as possible from as few explanantia as possible.[8]

In the rest of this paper I will undertake the hitherto unaccomplished task of defending unificationism about metaphysical explanation (*Metaphysical Unificationism*).[9] The main motivation I will offer to take such a view seriously has to do with the conceptual links between explanation and understanding. Explanation has been typically understood as a hybrid metaphysical-epistemological concept: explanations track an objective, worldly order but are also the sorts of things that increase understanding, make the phenomena intelligible, satisfy our curiosity, etc.[10] By contrast, the metaphysical explanation literature had a lot to say about the worldly aspect of explanation but much less about its epistemological aspect. Most authors don't believe that the kind of explanation they have in mind has epistemic or psychological connotations,[11] but even those who do rarely spell these out in detail.[12] This is surprising, since there is broad consensus that a complete theory of *scientific* explanation should shed some light on the relation between explanation and understanding. I would speculate that this feature of the grounding literature is partially responsible for the remaining hostility toward this notion: many philosophers are baffled by what it could even *mean* that some kind of explanation is dissociated from the things (understanding, intelligibility, relief from puzzlement, etc.) we normally associate with explanation.[13]

I will henceforth take seriously the idea that metaphysical explanation is a kind of explanation in the same sense scientific explanation is, and that it should therefore bear substantive links to understanding. The rest of the paper will proceed as follows. In Sect. 2 I will give a first-pass statement of the unification view. In Sect. 3 I will raise a trivialization challenge for this statement, offer a refined formulation to address it, and along the way also introduce the conceptual tools

---

[7] Hempel and Oppenheim (1948). Wilsch (2015, 2016) has recently developed a hybrid account of metaphysical explanation that combines elements of the covering-law view with the Backing Model: he takes metaphysical explanations to be arguments that cite metaphysical laws, which are in turn characterized in terms of "construction operations".

[8] Friedman (1974) and Kitcher (1981, 1989). See also Strevens (2008) for a view that combines elements of the causal and the unification views.

[9] I sympathetically discussed unificationism about metaphysical explanation (Kovacs 2017: 2943–6), but without offering a detailed account. Rather than defending full-blown unificationism, Roski (2017) lays out and tentatively defends a unification *constraint* on grounding. His Bolzano-inspired notion of unification is quite different from the one I'm working with here. Bertrand (2018) argues that a certain subset of metaphysical explanations (metaphysical explanations by constraint, as he calls them) are "top-down" rather than "bottom-up" and (similarly to explanations by unification) have a generalizing aspect to them.

[10] See Ruben (1990: 2), Kim (1994: 52, 53), Lipton (2001: 43ff.), and many others.

[11] Correia and Schnieder (2012: 24), Jenkins (2013: §5), Raven (2013: 193, 2015: 326) and Skiles (2015: 719).

[12] Trogdon's work (2013) is a welcome exception to this general trend, although I don't agree with his substantive views on the connection between metaphysical explanation and understanding.

[13] See Thompson (2016: 397) and Maurin (2018) for similar complaints.

needed to distinguish metaphysical from scientific explanation. In Sect. 4, I will spell out some of the view's more interesting consequences and respond to a number of objections. In Sect. 5, I will conclude that Metaphysical Unificationism offers an overlooked but promising account of metaphysical explanation.

2 Metaphysical Unificationism: a first pass

Unificationism belongs to a long-standing tradition of *argument views*: explanations correspond to arguments whose conclusion is the explanandum and whose premises are the explanantia. To keep things simple, I will represent arguments as Fitch-style trees (with the inference rules specified in a separate set of "filling instructions"—see below). The main advantage of this is that Fitch-style trees are reasonably close to mirroring natural reasoning and thereby help achieve the desideratum that metaphysical explanations increase understanding. There are other ways of representing arguments, but I have nothing to say about them in this paper.

Before launching into the substantive account, it's worth saying explicitly how I think of explanations and explanatory arguments. I prefer a view on which explanation is a relation between facts. Facts, as I think of them, are worldly entities made up of constituent individuals, properties and/or relations. As I think of them, they are "medium-fine-grained": necessarily co-extensional facts might not be identical, but different sentences may well describe the same fact. This becomes important when explanations are represented as arguments. Every sentence in an explanatory argument describes a fact, but different sentences can describe the same fact. Therefore, for our present purposes arguments with premises and conclusions that pairwise describe the same fact will count as identical. (Thus strictly speaking it might be more accurate to speak of equivalence classes of arguments rather than arguments. Since the substitution of a premise in a valid argument with a premise describing the same fact doesn't automatically preserve validity, I restrict my attention to those members of an equivalence class of arguments that are deductively valid.) While it is possible to be a unificationist without buying into this particular way of thinking about explanation and explanatory arguments (in which case specific details of the account I will present would need to be changed), I will presuppose this framework in the rest of the paper.[14]

Now back to unification. The intuitive idea behind unification theories is that explanation is a holistic matter. We can illustrate this by taking an ordinary causal-explanatory statement, for example, 'The window shattered because Jim threw a stone at it'. According to unification accounts, the truth of this statement is in part a function of how well it fits into a comprehensive system of similar statements. So, in wondering whether the stone throwing explains the window shattering, the unificationist may ask whether the shattering of vases, mirrors and plates could also be explained by stone throwings or whether replacing reference to stones with

---

[14] This framework also has the advantage of making it easier to state the naturalness criterion of unification; see Sect. 3 for details. Thanks to a referee for pressing me to get clearer on the relata of explanation and how they bear on explanatory arguments.

heavy objects would help us derive more cases of shattering. The more we can derive and the less we need to presuppose, the more unified our theory is; and the more unified the theory is that includes the throwing-shattering story, the more likely that story qualifies as an explanation.

The difficult task is to turn this impressionistic sketch into a reasonably clear *theory* of explanation. Nobody has done as much as Kitcher to accomplish this goal in the case of scientific explanation, and I will often rely on his account when developing my own. As we will see, many details (which I will discuss in due course) have to be revised to meet the demands of Metaphysical Unificationism. Let's start with some definitions (see Kitcher 1981, 1989 for similar ones).

A **schematic sentence** is the result of replacing some, though not necessarily all, non-logical expressions in a sentence with dummy letters (for example 'Grass is F', 'a is green', and 'a is F' are all schematic sentences).

A set of **filling instructions** for a schematic sentence is a set of instructions telling us how each dummy letter in a schematic sentence can be replaced.

A **schematic argument** is a sequence of schematic sentences.

A **classification** for a schematic argument is a set of sentences describing the inferential structure of that argument: it tells us which schematic sentences are premises, which ones should be inferred from other premises, and what rules of inference should be used.

An **argument pattern** is an ordered triple consisting of a schematic argument, the set of the sets of filling instructions of each term in the argument, and a classification.

Since my main focus is metaphysical explanation, I will draw my examples from metaphysics. Suppose, as is often thought, that an individual's instantiating a determinable property is explained by that individual's having a determinate of that determinable. Then we can give the following schematic argument:

Determinates and Determinables
(D$_1$) $\Phi(\tau)$
(D$_2$) If $\Phi(\tau)$ then $\Theta(\tau)$
(D$_3$) So, $\Theta(\tau)$

Our filling instructions will say that τ stands for a singular term, whereas Φ and Θ stand for predicates that refer to properties standing in the determinate-determinable relation. According to the classification, D_1 and D_2 are premises and the rule of inference used is modus ponens. Our argument pattern will then be the ordered triple of the schematic argument above, the set of sets of filling instructions describing the substitution instances of τ, Φ and Θ, and the classification. Most patterns will be much more complicated than this, but this simple example should suffice to get the main idea through.

Arguments are explanatory when they instantiate patterns that are members of a unified total set of argument patterns. We can now say a little more about what this means. Call the set of sentences that express the truths to be systematized 'K' and a

set of argument patterns that contain each member of K as a premise or as a conclusion a **systematization** of K. The core idea of unification is that a putative explanation is genuinely explanatory iff it's part of an explanatory theory, and that a theory is explanatory iff it's the best systematization of the truths.[15]

To clarify what makes a systematization good, Kitcher introduces the notion of *stringency*. The stringency of a set of argument patterns is a matter of how much similarity the patterns' logical structure shows and how hard it is to replace the non-logical expressions in each pattern. More precisely, stringency has a logical and a non-logical dimension. On the logical dimension, the more similar the argument patterns are to one another with respect to their logical structure, the more stringent the set of patterns they belong to is. For example, a set consisting of *Determinates and Determinables* and some pattern whose only permitted connective is the material conditional is more stringent than a set consisting of *Determinates and Determinables* and a pattern that permits other logical connectives, say, conjunction.

On the non-logical dimension, the more demanding constraints the filling instructions place on the non-logical expressions in the substitution instances of a pattern, the more stringent that pattern is. For example, compare *Determinates and Determinables* to

> *Unhelpful*
> (U$_1$) Δ
> (U$_2$) If Δ, then $\Theta(\tau)$
> (U$_3$) So, $\Theta(\tau)$

Let Δ be any sentence, τ any singular term, and Θ any predicate that stands for a determinable property. *Unhelpful* is much less stringent than Determinates and Determinables. First, it has less syntactic structure: Δ can be a sentence with *any* logical form, not just an atomic sentence assigning a one-place predicate to a singular term. Second, there are fewer non-logical constraints on (U$_1$) and (U$_2$) than on (D$_1$) and (D$_2$): Δ could be about anything whatsoever, not just about some individual instantiating a determinate property.

We can now give a more informative criterion of unification. For any range of phenomena to be systematized, the unifying power of a systematization of them is *directly proportional* to the overall stringency of its argument patterns (which in turn is directly proportional to the stringency of and similarity among its individual argument patterns) and *inversely proportional* to the number of argument patterns in it.

---

[15] The two 'iff'-s imply that if a theory is less unified than any other theory, it isn't explanatory at all. Woodward objects to this feature of Scientific Unificationism: surely a theory can be *somewhat* explanatory even if it's not *the* most unified theory (Woodward 2003: 367–369). Readers moved by Woodward's argument can distinguish between an all-or-nothing and a graded notion of explanation. While my account officially targets the all-or-nothing notion, it's easily extended to cover the graded notion: for any argument pattern, the better that pattern does on the criteria of unification, the more explanatory the arguments in it are.

As we will shortly see, this formulation needs more fine-tuning to avoid certain technical problems. Still, we already have enough details to see how the view would work in particular cases. In this regard, we face an immediate difficulty. In the scientific explanation literature, theories of scientific explanation are frequently tested against practice: it speaks against a theory if its implications radically diverge from the practices of working scientists. This kind of testing is more difficult in metaphysics, where consensus about the direction of metaphysical explanation is the exception rather than the rule.[16] To give a sense of how the view works, I will discuss two examples below: one that is uncontroversial and one where at least the source of the controversy is relatively uncontroversial.

First, most of us have the intuition that {Socrates} exists because Socrates does, rather than the other way round. Why? On the unification account, we need to evaluate both hypotheses against the backdrop of the best general argument pattern that accommodates them (to keep the phenomena to be systematized fixed, we will focus on the direction of explanation between *non-empty* sets and their members). Let's see, then, what an argument pattern would look like that allows us to explain the existence of {Socrates} by reference to the existence of Socrates. Presumably, it would look something like this:

> *Members to sets*
> (M_1) $\alpha_1...\alpha_n$ exist
> (M_2) If $\alpha_1...\alpha_n$ exist, then $\{\alpha_1...\alpha_n\}$ exists
> (M_3) So, $\{\alpha_1...\alpha_n\}$ exists,

where $\alpha_1...\alpha_n$ are any objects and n is finite.[17] Since any finite number of things have a set, the existence of any such (non-empty) set can be derived using some instance of *Members to sets*. (Note that in most cases, M_1 will itself be subject to a similar explanation, as well as *that* explanation's first premise, until we reach an instance of *Members to Sets* in which each α_i is an Ur-element or the null set.)

But what general argument pattern could accommodate the explanation of Socrates' existence by {Socrates}'? Here's a first stab:

> *Sets to members*
> (M_3) So, $\alpha_1...\alpha_n$ exists
> (M_2*) If $\{\alpha_1...\alpha_n\}$ exists, then $\alpha_1...\alpha_n$ exist
> (M_1) $\{\alpha_1...\alpha_n\}$ exists

But *Sets to members* gives rise to a dilemma. When applied with full generality, the pattern leads to an infinite regress: the existence of the Ur-elements and the null set is explained by the existence of the sets that could be formed from them, whose existence is in turn explained by the existence of *their* sets, and so on ad infinitum. Note that I'm not assuming anything like an absolute ban on infinite explanatory regresses here; all I help myself to is the modest assumption that if one

---

[16] As Lange (2014: 490) notes, there is a similar obstacle to assessing competing theories of mathematical explanation.

[17] Alternatively, we could lift the requirement that n be finite and introduce other restrictions to avoid Russell's paradox. I will forgo this task here.

systematization of the same range of phenomena leads to an infinite explanatory regress while another doesn't, then other things equal we should prefer the latter systematization. That is, *Members to sets* (understood unrestrictedly) should be preferred over *Sets to members* (understood unrestrictedly).

This leads to the second horn of the dilemma, according to which we need to put restrictions on *Sets to members*. Perhaps the pattern appropriately captures the direction of explanation between sets and their members as long as the latter are Ur-elements or the empty set, but the direction of explanation is swapped higher up in the set-theoretic hierarchy. So, perhaps the existence of {Socrates} explains the existence of both Socrates and {{Socrates}}. The problem with this view (its arbitrariness aside) is that it requires at least two separate argument patterns to systematize the target phenomena where one (*Members to sets*) would be enough, thereby making our argument patterns less unified. In short: if we want to avoid an infinite explanatory regress, then explaining Socrates' existence in terms of {Socrates}' commits us to a less unified systematization than explaining {Socrates}' in terms of Socrates'.

My second example is from material-object metaphysics. Facts about wholes are often thought to be explained by facts about their parts.[18] Schaffer (2010), however, defends priority monism, the view that the only fundamental material object is the cosmos, which is explanatorily prior to other objects. Schaffer's arguments for this view deserve a more thorough discussion than I have room for here, but two considerations stand out. First, he thinks that the existence of gunk would support priority monism, roughly because parts-to-wholes explanations could never bottom out if some of the parts were infinitely divisible. On the other hand, wholes-to-parts explanations do bottom out, given that the world isn't "junky" (not every object is a proper part of some other object).[19] The unificationist can agree that if the world is gunky but not junky, our best systematization is unlikely to be one that derives every fact about every object from facts about their parts and will more plausibly start with premises about the cosmos.[20] Another argument for priority monism is based on the assumption that the cosmos has emergent properties (properties not "fixed" by its parts), but that its parts don't have "submergent" properties (properties not "fixed" by the cosmos). Here, too, the unificationist has a good story about why this is relevant: if Schaffer's assumption is correct, the parts instantiating such and such properties and relations can be derived from the cosmos having

---

[18] See, for instance, Kim (1994: 67) and Skiles (2015).

[19] Schaffer (2010: 61–65).

[20] There is more to be said here. For example, perhaps the best set of explanations neither starts with the cosmos nor proceeds strictly to wholes from their parts but ultimately explains some facts about certain objects in terms of facts about those very objects (Raven 2016). Moreover, as a referee points out, if there is neither gunk nor junk then neither direction of explanation is supported solely on the basis of mereological considerations (though there may be several other criteria, such as naturalness, on the basis of which to declare a winner—see Sect. 3). My point is that we can make sense of Schaffer's gunk argument against the backdrop of Metaphysical Unificationism, not that the argument settles the debate.

certain properties, but not vice versa. So, the best systematization cannot start with the cosmos's parts to derive all the truths about the cosmos.[21]

Now, I won't pretend that Metaphysical Unificationism will replicate all the familiar views and arguments popular among grounding theorists. For example, as we will later see it's unlikely to license a number of principles thought to belong to the "impure logic of grounding". I will discuss these issues in Sect. 4.4, but we still have a long way to go before we get there. We first need to address a problem for my cursory statement of Metaphysical Unificationism. The tools I will use to address it will also be useful in distinguishing metaphysical explanation from scientific explanation.

3 Metaphysical Unificationism: refining the account

We aren't done yet; we need to address a technical problem that threatens to trivialize the account. The problem, which we can call the *Problem of Apparent Stringency*, is best seen as a special instance of a worry that any adequate theory of scientific explanation has to address. Take the following schematic argument:

One for All
(O_1) Σ
(O_2) So, Γ

Suppose Σ conjoins all the truths to be systematized. Since Γ is one of these truths, it's trivially derivable from Σ. In a sense, this pattern is quite stringent: the filing instructions are maximally specific (Σ has exactly one permitted substitution instance), and *One for All* allows us to derive everything else. If *One for All* qualifies as a unified theory, we would probably have to take that as a *reductio* of Metaphysical Unificationism. But does it?

Fortunately not, for Σ is extremely gerrymandered. After all, it's a lengthy conjunction of *all* the truths to be systematized, and while merely being a conjunction doesn't by itself make something gerrymandered, being a very long conjunction of truths (many of which are already gerrymandered) does.[22] Kitcher doesn't say much to cash out this intuition other than demanding that only projectable predicates occur in the premises of a scientific explanation.[23] I prefer a solution that relies on the notion of *naturalness*. Lewis's locus classicus (1983) is a

---

[21] Schaffer (2010: 57). As above, my purpose here is to show that unificationists can make sense of the argument, not that the argument actually works. Note that since Schaffer thinks of priority monism as a noncontingent thesis, his original arguments revolve around the *possibility* of gunk/junk and the *possibility* of submergence/emergence. I rephrased these arguments because from a unificationist perspective, there is no reason to assume that the direction of explanation is invariant across possible worlds (cf. Jones 1995b).

[22] We could replace reference to all truths with reference to all facts (i.e. we could demand that no two conjuncts in Σ describe the same fact) or perhaps all atomic facts; I don't think these tweaks will make much of a difference to the argument, as these too strike me as highly gerrymandered. Thanks to a referee here.

[23] Kitcher (1989: 488–489).

useful starting point: he characterized (but declined to define) natural properties as empirically discoverable properties that track similarity in nature, are preserved between duplicates, figure in laws, feature in a minimal supervenience base and are intrinsically eligible to be the semantic values of our predicates.[24] Other plausible theoretical roles include the conferring of causal powers.[25]

For technical reasons that I saw best to relegate to a footnote, in addressing the Problem of Apparent Stringency I appeal to the notion of a natural *fact* rather than that of a natural property.[26] Natural facts are not entirely unheard of in metaphysics and have already been adopted to do work similar to that of natural properties[27]; more generally, Sider (2011: Ch. 6) has argued convincingly that nothing in Lewis's notion prevents us from applying it to entities other than properties. (Of course, it's already standard to talk about *fundamental* facts. But as I will make clear in Sect. 4.1, by 'natural fact' I don't mean what most authors mean by 'fundamental fact'.) Officially, I will treat naturalness as coming in degrees; that is, I won't presuppose that any fact is perfectly natural.

A further departure from Lewis's original notion is that I recognize *two* graded notions of fact-naturalness, which I will call metaphysical and scientific naturalness, or in short M-naturalness and S-naturalness. (The distinction will serve a dual role in both solving the Problem of Apparent Stringency and clarifying the distinction between metaphysical and scientific explanation.) A similar distinction can be found in Dorr and Hawthorne's work (2013), who note that we might treat certain distinctively metaphysical notions, or as they put it structural properties and relations (parthood, compresence, instantiation, etc.), as in an important sense natural despite doing poorly on some of Lewis's original criteria (for example they don't figure in the laws of nature and aren't empirically discoverable). By contrast, there may be properties that are crucial for scientific theorizing but don't do well on Lewis's more metaphysical criteria (e.g. preservation between duplicates or being part of a minimal supervenience base).[28]

Analogously, I will distinguish between M-natural and S-natural *facts*. Similarly to property-naturalness, these notions resist analysis in simpler terms but can still be characterized informatively. Examples come easily: that such and such charge and

---

[24] See also Armstrong (1978) for an alternative development in terms of universals.

[25] See Schaffer (2004).

[26] If we formulate the non-gerrymandering constraint in terms of properties, then given the requirement that explanatory arguments be deductively valid it will follow that (apart from a few special cases, e.g. logical truths, which we can safely ignore here) any predicate that occurs in the explananda also has to occur somewhere in the explanantia. But then, any complete systematization contains the same set of predicates in the premises (namely, all of them) and *a fortiori* a set of predicates that stand for overall equally natural properties. We can avoid this problem by replacing natural properties with natural facts in the account. (Thanks to Louis deRosset and Cian Dorr for pressing me on this problem.).

[27] See Hirsch (2008: 523) and Dunaway (ms).

[28] In the same paper, Dorr and Hawthorne also argue that some of the roles in question might not be jointly satisfiable by any candidate notion of naturalness. Unfortunately, I lack the space to adequately address this important challenge. All I can to do is state, without argument, my preferred way of dealing with it: S-natural and M-natural properties are those that, on the whole, best realize a weighted total of the roles associated with the relevant notion of naturalness.

mass properties bundle (via compresence) into an electron is a good candidate for being an M-natural fact, and that the Earth orbits around its and the Sun's barycenter is a good candidate for being an S-natural fact. Moreover, as Sider (2011: Ch. 6) notes, some of the theoretical roles Lewis assigns to natural properties have reasonably clear applications to facts. For example, S-natural facts tend to be more similar to the laws of nature and their instances than non-S-natural facts are, and they concern causally efficacious individuals with unified causal profiles.[29] We can also formulate fact-friendly versions of other roles: the M-natural facts jointly constitute a complete supervenience base for all the other facts, and S-natural facts are empirically discoverable. Moreover, other things equal the similarity between S-natural and M-natural facts that involve the same individuals or the same properties and relations will be greater than between non-natural facts with similarly "swapped" constituents. Finally, S-natural as well as M-natural facts can be intrinsically eligible to be expressed by our sentences the same way natural properties are intrinsically eligible to be expressed by our predicates.[30]

The above is by no means meant to be an exhaustive characterization of S-natural and M-natural facts, and each of the features mentioned above is to some extent negotiable. Ultimately I take these notions to be primitive, but I hope to have said enough to make it plausible that we have a solid grasp on them. We can now plug M-naturalness and S-naturalness into Kitcher's anti-gerrymandering constraint. When choosing metaphysically explanatory argument patterns, one criterion of their unification is the extent to which the sentences that serve as substitution instances in them correspond to overall M-natural facts. Ditto for scientific explanation, except that this time we should favor patterns with sentences corresponding to S-natural facts. We can now also see why *One for All* is unexplanatory: its single premise is an incredibly long and complicated sentence, plausibly corresponding to an extremely unnatural fact (both in the scientific and in the metaphysical sense of naturalness).

It's worth drawing attention to a few more salient features of the account presented here. One is that naturalness is a graded constraint: argument patterns that display an overall higher degree of naturalness are better, but I don't require the

---

[29] Causal powers are usually assigned to properties, but Baysan (2018) argues, convincingly in my view, that this shouldn't be understood literally and that the ultimate bearers of causal powers are individuals. If he is right, then facts are at least as well suited for locating causal powers as properties are.

[30] This last role raises tricky questions about the relation between the intrinsic eligibility of a fact to be expressed by a sentence and the intrinsic eligibility of a property to be expressed by a predicate that occurs in that sentence. As a referee observes, allowing the two to come apart might violate plausible principles of compositionality. While I have no settled view on the relation between the eligibility of facts and of properties to serve as semantic values, I think we can preserve compositionality without committing ourselves to any particular view. This is because intrinsic eligibility is just one factor that determines naturalness. Perhaps the intrinsic eligibility of a fact to be expressed by a sentence is a function of the intrinsic eligibility of certain individuals, properties and relations to be expressed by the sentence's constituents. But it doesn't automatically follow that the M-naturalness of a fact is likewise a function of the M-naturalness of the respective individuals, properties and relations, since even if meaning-eligibility obeys principles of compositionality, other factors that go into the weighed total that ultimately determines an entity's naturalness might not (cf. footnote 28).

patterns to only feature *perfectly* M-natural facts. Given some plausible ancillary assumptions, such a requirement would likely leave us with an insufficiently rich array of explanans facts from which to derive the explananda. But not imposing the requirement is also plausible for independent reasons. Even if every sentence were derivable exclusively from sentences describing perfectly natural facts, it shouldn't automatically follow that the resulting set of derivations would correspond to the best systematization: perhaps some phenomena are on the whole best unified by a systematization that includes less than perfectly M-natural facts.[31] Moreover, it's not even obvious that there *are* perfectly M-natural facts in the first place.[32] By severing explanatory basicness from perfect naturalness, we can leave this possibility open without having to take on board infinite explanatory regresses.[33]

Another interesting upshot of my view is that some explanations might qualify as both scientific and metaphysical. If an argument instantiates a pattern in the best systematization according to both the scientific and the metaphysical criteria (because it features facts that do well on both the S-naturalness and the M-naturalness dimensions), it will qualify as a scientific and as a metaphysical explanation at the same time. I think this is a good result: some scientific explanations, e.g. those that explain an object's properties in terms of its microstructure, have all the characteristics typically associated with metaphysical explanations. There is no good reason to think of scientific and metaphysical explanation as mutually exclusive categories.

Let me restate the unification account of metaphysical explanation, this time with the naturalness constraint included. The facts that $A_1 \ldots A_n$ metaphysically explain the fact that B just in case the argument whose premises are $A_1 \ldots A_n$ and whose conclusion is B belongs to the most unified set of argument patterns. The unifying power of a set of argument patterns is (1) *directly proportional* to the stringency of the total set of argument patterns, as well as of the individual argument patterns within them; and it's *inversely proportional* to (2a) the number of argument patterns, (2b) the number of premises in them, and (2c) the overall M-naturalness of the facts those premises describe. (The facts that $A_1 \ldots A_n$ scientifically explain the fact that B under the same conditions, except that we replace M-naturalness in (2c) with S-naturalness.)

This is a complicated and messy account, especially since I have said nothing about the relative weights and ranking of these criteria. In some special instances they are precisely comparable,[34] and in many others we can reach a verdict without an exact ranking. But there doubtless remain hard cases we won't know how to approach. However, I don't think we should be too worried about this. Once we are on board with the idea that metaphysical explanation is a holistic matter, we should

---

[31] Cf. Jones (1995a).

[32] Cf. Armstrong (1978: 67–68), Sider (1993: Ch. 3.22), and Williams (2007) for a similar possibility about natural properties.

[33] We might even use the distinction to give an error-theoretic account of why such regresses might *seem* possible: we are prone to mistaking the possibility that naturalness has no lower bound for the possibility that explanation isn't well-founded. See also Sect. 4.1.

[34] Cf. Kitcher (1989: 488–494).

recognize that theory comparison *is* often messy and complicated. In this regard, metaphysics isn't radically different from science, and we shouldn't expect from a theory of metaphysical explanation any less complexity, or more reliable practical guidance, than we would from a theory of scientific explanation.

This concludes my presentation of Metaphysical Unificationism. In the next section, I will highlight some interesting connections my account brings out between metaphysical explanation and other notions, as well as respond to a number of worries concerning the account.

4 Some consequences of Metaphysical Unification

In this section I will spell out some further consequences of Metaphysical Unificationism and consider a few objections to the account. In Sects. 4.1–4.3 I will discuss the view's implications for the relation between fundamentality and naturalness, metaphysical explanation and understanding, and the possibility of explanatory indeterminacies. In Sects. 4.4–4.6 I will discuss three problem cases: logical grounding, the explanation of logical truths, and the explanatory status of modal truths. While Metaphysical Unificationism has controversial consequences for all of these issues, I will argue that none of them is unacceptable.

4.1 Fundamentality and naturalness

According to a standard definition, a truth is fundamental just in case it's metaphysically unexplained ("independent", to use Bennett's expression).[35] As Bennett (2017: 127–128) points out, fundamentality in this sense and naturalness are distinct notions: something could be independent despite manifestly failing on Lewis's criteria of naturalness, and conversely, we can imagine perfectly natural entities that aren't independent. For example protons might be non-fundamental but perfectly natural, whereas facts about indivisible but psychologically sophisticated souls might be fundamental but not perfectly natural (the first example is Bennett's, while the second is mine). But even those who agree with Bennett on this might feel that there ought to be *some* connection between fundamentality and M-naturalness. And indeed, on my unificationist view there is an important (if indirect) connection between them.

To see this, notice first that the unificationist is free to adopt the independence definition of fundamentality. It's just that what's fundamental will be settled holistically: a fact is fundamental just in case it's described by a sentence among the permitted substitution instances of a schematic sentence that occurs as a premise in the most metaphysically unified set of argument patterns. Next, recall the criteria specified in Sects. 2 and 3 for finding the most unified set of argument patterns. One of these was that the more M-natural the facts are that the sentences in a set of

---

[35] See also Schaffer (2009: 373, 2010: 38), deRosset (2013: 5), and Skiles (2015: 726). Bennett herself understands independence in terms of a category-neutral notion of *building*, rather than grounding or metaphysical explanation.

argument patterns describe, the more unified that set is (other things being equal). For this reason, we can expect that other things equal the fundamental facts will be more M-natural than the non-fundamental ones. To be sure, this is a rather tenuous connection between fundamentality and M-naturalness, but it's an important connection nonetheless.

4.2 Explanation and understanding

In the introduction I said that a good theory of metaphysical explanation ought to shed light on how metaphysical explanations increase understanding. Now, in one sense any theory of metaphysical explanation is at least *consistent* with the requirement that explanations increase understanding. After all, theorists of all stripes can insist that explanations are answers to why-questions. (They don't need to be *solely* answers to why-questions; it's sufficient if all explanations have the property of being the answer to a why-question). Moreover, in so far as explanations answer why-questions, they cannot lack the potential to increase understanding.[36]

Thus, the problem with backing views isn't that they are incompatible with the understanding requirement. Rather, on these views it's not transparent what it is about backing relations that yields understanding or, if you prefer, answers to why-questions. Suppose the grass's being of shade $green_{17}$ *grounds* the grass's being green. It isn't obvious how the presence of grounding answers the question of why grass is green or how it helps us understand why grass is green. Presumably it's not an analytic truth that if grass's being of shade $green_{17}$ grounds grass's being green then the former fact also metaphysically explains the latter. But then it's unclear how much the introduction of grounding adds to the plausible but fairly generic point that explanations answer why-questions.[37]

One attraction of Metaphysical Unificationism is that allows us to say something more specific about *how* explanations increase our understanding. Metaphysical explanations provide a fundamentally global sort of understanding: they help us see how a large number of phenomena are the consequences of a small number of basic facts, from which they can be derived using relatively few and similar patterns of derivation. By identifying many similar derivations, we start noticing connections and common patterns that tie the phenomena together. The enfolding web of connections, all branching from a few basic premises, enhances our understanding.[38]

That explanation increases understanding doesn't mean that it always increases everyone's understanding. Not every act of communicating an explanatory product is an explanatory act. For example, the explanation may be old news, or the audience may lack the conceptual resources to understand it.[39] Rather, genuine explanations should have the potential to increase the understanding of rational

---

[36] Thanks to an anonymous referee for pressing this point.

[37] Cf. Maurin (2018).

[38] Cf. Friedman (1974: 18–19), Kitcher (1981: 529–530) and Kitcher (1989: 431–432).

[39] Lewis (1986b: 226–228).

agents with proper training (e.g., someone who has an adequate grasp of the literature on properties, facts, sets, and other metaphysical issues). Metaphysical Unificationism accounts for why metaphysical explanations have this potential.

To link explanation to the understanding of properly trained rational agents is reminiscent of a suggestion Michael Friedman made long ago in his now classic paper on unification. As Friedman points out, just because explanation is a partly epistemic/psychological notion, it doesn't follow that it's in the eye of the beholder: explanatoriness may be in an important sense objective if "what is scientifically comprehensible is constant for a relatively large class of people" (1974: 7–8). I suspect that a failure to appreciate the distinction between explanation being psychological and its being agent-relative is at least in part responsible for the contemporary literature's neglect of the role metaphysical explanation plays in producing understanding. Once we realize that the notion can be objective *and* epistemic/psychological, we should no longer feel uncomfortable with the idea that metaphysical explanation has one leg in the metaphysical and another in the epistemic realm.

4.3 Explanatory indeterminacy

To metaphysically explain something is to find its place in the metaphysically best systematization. But what if there is no such thing as *the* best systematization? One might find indeterminacy about metaphysical explanation hard to swallow. How can it be indeterminate what is prior to what or which facts are fundamental? I won't try to convince skeptics that explanatory indeterminacies are acceptable in metaphysics. Instead, I will mention two factors that, to my mind, mitigate their counterintuitiveness.

First, while nothing in Metaphysical Unificationism rules out ties, they are extremely unlikely to occur in our world. Unification has multiple dimensions that are difficult to quantify. While this doesn't rule out explanatory ties, lots of things would need to fall in place for one to actually occur. Moreover, in most cases when we wonder about the direction of metaphysical explanation, together with the first-order truths our criteria will yield a relatively clear answer. We can easily confirm this by focusing on cases in which there is broad agreement on these first-order truths, for example the Socrates/{Socrates} case.

Second, the indeterminacy problem seems less worrisome if we keep in mind what metaphysical explanation *is* according the Metaphysical Unificationism. It isn't the tracking of "backing" relations; nor does it have anything to do with essence or real definition. Instead, it's simply a matter of finding the best systematization of the facts according to the criteria specified at the end of Sect. 3. And even if it's implausible that there is any indeterminacy about the pattern of backing relations holding in our world, it isn't especially puzzling how it could be indeterminate which systematization scores highest on these criteria.[40]

---

[40] It's worth noting that there may well be other legitimate senses of 'fundamental' in which even if there's indeterminacy about what metaphysically explains what, certain things might nonetheless be fundamental. For example there could still be mereologically simple ("compositionally fundamental")

4.4 Unificationism and the impure logic of grounding

In the grounding literature it's standardly assumed that that there is an "impure logic" of metaphysical explanation: a set of formal principles guiding the interaction of metaphysical explanation with the logical constants. For example disjunctions are supposed to be explained by their true disjuncts, conjunctions by their conjuncts taken together, existentially quantified truths by their witnesses, and so on.[41] But on the present view, there is unlikely to be such a thing as the impure logic of metaphysical explanation. This is because for the unificationist, similarity in logical structure is just one of the features that make a set of argument patterns unified, and its absence may well be offset by other (non-formal) virtues, for example similarity in non-logical structure or the candidate explanantia's naturalness.

Here's an example of how this could happen. Take a view according to which properties "confer" causal powers to their bearers. There are several ways of thinking about properties that could motivate such a view. Perhaps the fact that some x has property F explains why x also has causal powers $P_1 \ldots P_n$; or perhaps x's having F *just is* x's having powers $P_1 \ldots P_n$. Either way, when such a view is combined with Unificationism, it will be plausible to explain why x has P_i (for any $i \in \{1 \ldots n\}$) by appealing to the conjunctive fact that x has $P_1 \ldots P_n$. If F has a sufficiently unified causal profile, the fact that x has P_1 & x has P_2 &…& x has P_n will be a natural candidate by which to explain why x also has P_i. After all, lots of facts about individuals bearing causal powers can be similarly derived by citing fairly natural facts about the complete causal profiles of those individuals, which will result in an overall fairly unified set of derivations. Notice, however, that in these cases the direction of explanation runs contrary to the supposedly general rule that conjunctions are explained by their conjuncts.[42]

I don't expect this example to be uncontroversial. The point is that once we allow similarity of logical structure to be outweighed by other considerations, it's overwhelmingly likely that at least in some cases it will be so oughtweighed. So, while it's in principle conceivable that some fully general formal principles occur as patterns in the best systematization, the complex set of criteria I proposed in Sect. 3 are unlikely to classify any of them as such; it's just hard to see how the best systematization could contain any argument pattern with *no* restriction on its non-logical expressions.

Footnote 40 continued

things, and as should be clear from Sect. 3, there could also be instantiations of perfectly M-natural properties.

[41] See, for instance, Correia (2010, 2014), Fine (2010) and Schnieder (2010).

[42] It is important to note that on the view I envisage, the plurality of facts that x has P_1, that x has P_2, …, and that x has P_n do *not* jointly explain why x has P_i. For assuming the plausible principle that any member of a plurality of facts that fully explain an explanandum partially explains that explanandum, it would then follow (absurdly) that x's having P_i partially explains itself. (See Kovacs 2018 for arguments that this result is indeed absurd).

The unificationist can respond to this worry in at least three ways. The most modest response is to simply restrict the scope of Metaphysical Unificationism and treat logical explanation as a sui generis category that requires separate treatment.[43] While defensible, I find this modest response suboptimal. To my mind, one of the main attractions of the unification account is that it can treat diverse forms explanation as instances of the same general phenomenon.

A more ambitious response is to modify Unificationism within the realm of logical grounding to bring it in line with the formal principles. Roski (2017) has recently proposed a kind of unification constraint on grounding, largely inspired by Bolzano's work, according to which we can derive a maximal number of conclusions from a minimal number of premises that ground them. A further constraint (whose main role is to prevent arguments similar to *One for All* from qualifying as genuine explanations) is that no truth can be grounded in any truth with more logical complexity than itself. The unificationist might identify logical explanation with unification in this broadly Bolzanian sense (though note that this goes beyond the Bolzano–Roski view, which uses unification merely as a constraint) and could argue that this modified notion of unification can be reconciled with, and perhaps even justify, the principles of the impure logic of grounding. While more attractive than the first solution, I don't find this response fully satisfactory either. As Roski himself notes, Bolzano's notion of unification is very different from Kitcher's (and consequently from the one employed here, which is based on it). This would allow us to understand metaphysical and scientific explanation as species of the same genus. Still, logical explanation (with no naturalness criterion) would turn out to be a wholly different kind of relation from metaphysical explanation (with no formal constraint on logical complexity). So, adopting the Bolzano–Roski view for logical explanation would still bifurcate the notion of explanation, albeit less so than the first strategy.

The third response is the most radical one, and the one I ultimately favor. It is to contend that there is no such thing as logical explanation and accordingly no such thing as the impure logic of grounding. I expect that this response would strike many as heavy-handed at best and unacceptable at worst. Not me. For my part, I never found the principles supposedly guiding the impure logic of grounding particularly compelling in the first place. In fact, I'm inclined to think that some of them have *no* true instances. Conjunctive grounding is a case in point: I can think of a number of acceptable responses to the question, "Why is it dark and rainy?", but "Because (i) it is dark, (ii) it is rainy" isn't one of them; I could hardly think of a less informative answer to it.[44]

---

[43] Wilsch (2016: 20–21) suggests a similar response to the objection that his D–N theory of metaphysical explanation cannot accommodate logical grounding.

[44] Marshall (2015: 3162, 3163) rejects the principle that universal generalizations are grounded in all of their instances taken together on the basis of parsimony considerations: the principle would force us to accept lots of particular facts as fundamental, which someone denying the principle could derive from a much smaller array of general facts. While Marshall neither endorses unificationism about explanation nor extends the reasoning to other formal principles, his reasoning bears an obvious similarity to mine. It's also worth noting that Marshall argues against the groundedness of universal generalizations in their

But I don't want to rely only on my intuitions here, which I expect many readers won't share. Even if the impure logic of grounding is somewhat intuitive, it's not a non-negotiable part of our notion of metaphysical explanation. There is no logical theorem each instance of which corresponds to a scientific explanation, and I see no obvious reason why we should think of metaphysical explanation in radically different terms. Influential as Fine's formally oriented work on metaphysical explanation is, he never offered independent motivation for the impure logic so widely presupposed in the contemporary grounding/explanation literature. If an otherwise attractive theory of metaphysical explanation doesn't license the rules of the impure logic of grounding, I propose that we scrap those rules rather than the theory. To readers sympathetic to my account who are unwilling to go quite this far, I recommend one of the more moderate options: treat logical grounding either in non-unificationist terms or (preferably) along the lines of the Bolzano–Roski view.

4.5 The explanation of logical truths

In the previous section I discussed putative metaphysical explanations that are also logical truths. I now want to turn to a related but importantly different issue: metaphysical explanations whose *explanantia* are logical truths.

For the metaphysical unificationist, the explanatory status of logical truths is a matter of which derivations with logical truths as their conclusion belong to the best systematization. This quickly leads to a puzzle. According to orthodox grounding-theoretic approaches, at least some logical truths aren't fundamental. For instance, according to the standardly accepted impure logic of grounding, disjunctions are grounded in (and so explained by) their true disjuncts, which implies that any instance of $P \vee \sim P$ is explained by either P or $\sim$ P. Above I suggested dispensing with the impure logic of grounding, but even then it wouldn't follow that *no* logical truth has a metaphysical explanation.

But if metaphysical explanation involves derivation from premises, there is a case to be made that *no* logical truth has an explanation. This is because it's at least somewhat plausible that our best systematization will derive all logical truths from zero premises. After all, whenever a systematization derives some logical truths from premises that stand for some not-perfectly-natural facts that show a certain degree of similarity, there will always be a zero-premise derivation that shows even more cohesion and which takes nothing away from the overall naturalness of our patterns.[45] Both claims are true, since zero-premise derivations rely on exactly the same explanans facts (none), which in turn have same amount of less-than-perfect naturalness (again, none).

Footnote 44 continued

instances in the context of defending a Humean account of laws, which is structurally very similar to unificationist accounts of explanation.

[45] Of course, this presupposes a controversial feature of classical logic, namely that tautologies follow from zero premises. I will proceed on this assumption, since it makes my job harder; if zero-premise arguments are invalid, unificationists face no special challenge about the explanatory status of logical truths.

I'm tempted to see this as a feature of my approach rather than a bug. There *is* something intuitive about the notion that logical truths have no explanation. By saying so, my account might even help make sense of the commonly heard but vague slogan that logic is "not about the world" without implying an objectionable sort of logical conventionalism.[46] This being said, Metaphysical Unificionism has the necessary tools to satisfy those who don't share my intuition, or at least to meet them halfway. Earlier I defined fundamentality as metaphysical unexplainedness, but as the possibility of explanations that correspond to zero-premise arguments reveals, 'unexplained' is ambiguous between *not being an explanandum* and *not having an explanans*. These are distinct notions. If C isn't an explanandum, it's not the conclusion of any derivation in the best systematization; while if it merely doesn't have an explanans it might still be the conclusion of a derivation, albeit a zero-premise one. Call the former notion *strong* and the latter *weak* unexplainedness.[47] The distinction between the two forces us to choose between two disambiguations. If fundamentality is weak unexplainedness then plausibly every logical truth is fundamental, since the best systematization doesn't derive any logical truth from more than zero premises. But if fundamentality is strong unexplainedness, then for the same reason no logical truth is fundamental, since none has to be assumed as a premise in the best systematization.[48]

The issue is largely semantic: how should we use the word 'fundamental'? In my view, strong unexplainedness yields a notion of fundamentality closer to our intuitive conception, since it allows us to maintain that there is a sense in which some logical truths are more fundamental than others. Suppose T_1 and T_2 are logical truths, and in the best systematization T_2 is derived from T_1, which serves as an intermediate conclusion in a Fitch-style tree with zero premises. What should we say, then, about T_1's and T_2's fundamentality status? If both are absolutely fundamental, then neither can be more fundamental than the other. But if neither is absolutely fundamental, it's possible for the one to be more fundamental than the other, consistently with our usual understanding of absolute fundamentality.

Now, it would be misleading to say that T_1 explains T_2; both T_1 and T_2 are ultimately derived from zero premises, so neither has an explanans (although each is explained). Yet, there is a sense in which T_1 is explanatorily prior to T_2: it comes before T_2 as a step in the argument whose conclusion is T_2 and which is T_2's explanation. To be clear, I'm not *arguing* that the intermediate conclusions of an argument are more fundamental than those of its conclusions that occur later. Rather, I'm trying to *stipulate* a notion of relative fundamentality that tracks a certain kind of priority. This helps us keep the intuition that in some sense not all the

---

[46] Cf. Sider (2011: 97–104).

[47] Weak unexplainedness is similar to Fine's (2012) notion of zero-ground: being ungrounded isn't the same thing as being grounded by a set of grounds that is identical to the empty set. Of course, Fine himself rejects the idea that all logical truths are zero-grounded.

[48] Of course, we often need to appeal to inference rules when we give explanatory arguments for logical truths. However, these inferences are not themselves among the explanantia of these truths. Here I agree with deRosset (2013: 20–21), who argues that not everything that appears in the explanation of a certain explanandum is among the explanantia of that explanandum.

logical truths have the same explanatory and fundamentality status even if none of them is absolutely fundamental.

The conclusion that all logical truths are unexplained and non-fundamental might strike some as unacceptable. But I hope to have shown that it's more palatable than it seems at first glance. For one, we can distinguish two senses of unexplainedness, not being an explanandum and not having an explanans, and point out that logical truths are unexplained only in the latter sense. For another, we can maintain that some logical truths are explanatorily prior to and more fundamental than others at least in the sense that they occur earlier in the derivations that belong to our best systematization.

4.6 Fundamental modality?

Until now, all my examples were confined to truths that aren't prefixed with a modal operator. What about the explanatory status of modal truths? Some philosophers accept reductive accounts of modality: for example, Lewis (1986a) thinks that modal truths reduce to truths about the modal pluriverse, whereas Fine (1994) and Kment (2006) attempt to reduce them to essential truths. But most modal truths cannot be derived from any combination of non-modal ones. Doesn't this mean that at least some modal truths are fundamental? And doesn't that imply that unificationists are committed to modality at the fundamental level?

There are two ways of addressing the challenge. The first one, which stays closer to the way contemporary metaphysicians tend to think of reductionism about modality, is to amend the premises with "bridge laws" that connect the modal operators to the vocabulary of FOL (for example, 'There is a world such that A iff possibly A'). I expect that many will object to this strategy on the basis that it violates Sider's (2011) principle of Purity, according to which only fundamental concepts can appear in a fundamental truth: the bridge principle 'There is a world such that A iff possibly A' contains a modal concept and so, the thought goes, it cannot be fundamental; but then, how can it appear among the premises in our best systematization?

To evaluate this objection, I first need to recast it in my own terms. In my conceptual framework Purity disambiguates into three distinct theses, one of which is downright ill formed, the second is implausible, while the third may be true but has no bearing on reductionism about modality. Sider himself formulates Purity in terms of a general notion of fundamentality (or "structure", as he calls it), a category-neutral notion that applies to propositional as well as non-propositional aspects of reality. By contrast, I distinguish between fundamentality and M-naturalness. The former is reserved for facts, and I understand it in line with most of the grounding/explanation literature: a fact is fundamental just in case it is unexplained (weakly or strongly—see the previous sub-section). Understood in terms of this notion of fundamentality, Purity is ill formed since nothing that is not itself a fact has a fundamentality status.

This leads to the second and third interpretations. Unlike fundamentality (as I understand it), M-naturalness covers facts as well as properties (and perhaps other non-propositional aspects of reality). Earlier, I distinguished between a graded and

an absolute notion of M-naturalness. Since Purity is a thesis about absolute fundamentality, it is the latter that is of concern here. We can then ask two questions: (i) whether fundamental facts only involve perfectly M-natural concepts and (ii) whether perfectly M-natural facts only involve perfectly M-natural concepts. For reasons mentioned in Sect. 4.1, the answer to the first question is clearly 'No'; for the unificationist, part of the rationale to distinguish between fundamentality and M-naturalness is precisely the live possibility that the explanatorily most basic (i.e., fundamental) facts are not all perfectly M-natural. This is the sense in which Purity is well formed but false. By contrast, the answer to (ii) may be 'Yes'; one might even try to argue that we lack a firm handle on the notion of M-naturalness unless we accept it. However, since I deny that reductionism about modality requires the explanation of modal facts in terms of facts that only involve perfectly M-natural concepts, this disambiguation of Purity lacks bite even if it's true.

In short, then, the only thesis in the vicinity of Purity that a unificationist has reason to endorse is fully consistent with reductionism about modality. Still, some might feel that this answer doesn't quite get to the heart of the matter: the real issue is that a systematization that takes some modal truths for granted (i.e. "bridge laws" connecting the modal truths to their putative reduction base) doesn't give us a truly reductive theory of modality. Rejecting Purity doesn't make this intuition go away.[49]

I'm somewhat receptive to this objection, which leads to my second and preferred answer to the challenge: we shouldn't think of reductive theories of modality in terms of explanation at all. To be sure, this is how they are typically treated in the contemporary literature. But this is by no means mandatory; I think we should treat attempts to reduce modality as attempts at conceptual analysis. Some philosophers treat grounding and metaphysical explanation as being intimately connected to reduction; for example Rosen maintains that if p reduces to q and p is true, then the fact that q grounds the fact that p (2010: 122). I, by contrast, think that grounding and reduction are orthogonal notions (cf. Audi 2012). It's worth noting that in the context of the Rosen-Audi debate, 'reduction' has been taken to mean "real definition": roughly, a kind of definition that captures the essence of worldly items rather than concepts or linguistic expressions. That is, p reduces to q just in case for it to be the case that p *just is for it to be the case that* q, in the real definition sense of 'just is for it to be the case' (Rosen 2010: 123).

As I'm skeptical about non-modal notions essence as well as the accompanying notion of real definition,[50] this is emphatically *not* what I mean by 'reduction'. Instead, what I mean by it is a hybrid metaphysical-cum-epistemic notion (incidentally, also the notion usually meant by 'reduction' in much of the literature in the philosophy of mind and the philosophy of science)[51]: a special case of numerical identity. On this conception, if a reduces to b then (i) a is numerically

---

[49] Cf. Kim's (1998: 96) objection to the classic Nagel (1961) model of reductive explanation.

[50] My thinking about this topic has been largely influenced by Cowling (2013).

[51] See, e.g., Smart (1959), Crane (2001) and van Gulick (2001).

identical to b and yet (ii) it isn't the case that b reduces to a. How can (i) and (ii) both be true? The word 'reduces' creates opaque contexts, so the truth-value of 'reduces'-sentences is sensitive to the conceptual guises under which the putative reduction relata are presented.[52] What the conceptual guises need to be like for the respective 'reduction'-sentence to be true arguably cannot be specified in a topic-neutral way, but generally we should expect the expression on the right-hand side of a true 'reduces'-sentence to provide more information about the structure of the same entity than the expression on the left-hand side. This is why, for instance, 'water reduces to H_2O' is true whereas 'H_2O reduces to water' is false. In short: the *metaphysical* relation underlying reduction is plain numerical identity; yet reduction is an asymmetric relation due to the informational asymmetry between the conceptual guises under which the entity to be reduced is presented. So, reduction (as I understand it) is a kind of conceptual analysis.

Most purportedly reductive accounts of modality can be understood as attempts to "reduce" modality in the aforementioned sense. Lewis's modal realism is a case in point: his goal wasn't to provide a real definition of modal truths, or the essence of modality, or anything like that (indeed, Lewis was fairly suspicious of these notions). Rather, he tried to show that we could get rid of primitive modality by trading in the language of boxes and diamonds for the language of possible worlds and counterpart relations. From this perspective, there is no problem of explaining modal truths reductively. Properly understood, on reductive views modal truths aren't to be explained at all; they are to be analyzed, and then their analysans might be amenable to a metaphysical explanation. Thus understood, Metaphysical Unificationism is fully compatible with reductionism about modality.[53]

5 Concluding remarks

In this paper, I developed a unification account of metaphysical explanation. While I offered no knockdown argument for it, I made the case that it has several advantages. It gives us a way to make sense of metaphysical explanation without buying into the Backing Model. It clarifies the link between fundamentality and naturalness without objectionably identifying the two. It provides a unified account of metaphysical and scientific explanation, thereby making both Metaphysical and Scientific Unificationism more attractive. Finally, it ties metaphysical explanation

---

[52] Cf. van Riel (2013).

[53] It's not always entirely clear whether a purportedly reductionist theory of modality would survive when reinterpreted along these lines. For example, according to Sider's neo-Humean view, we classify certain truths as 'necessary' and others as 'contingent' as a matter of convention; though what's necessary isn't up to us, it's up to us how we use the word 'necessary', and our use is not superior to alternative, non-actual uses (Sider 2003, 2011: Ch. 12). 'Necessary truth' would then be defined as membership in an exhaustive list of categories ('mathematical truth', 'logical truth', etc.). It may well be an impossible task to specify all these categories, but Sider doesn't base his claim of having a reductive theory of modality on actually having specified them. Would his view qualify as reductionist when reinterpreted as a conceptual analysis of 'necessary truth'? In my view, without an actual list we only have the *recipe* for a reductive account, but not yet the account itself.

back to the notion of understanding, where it properly belongs. My brand of Metaphysical Unificationism also has some surprising features: nothing metaphysically explains the logical truths; it doesn't (at least in the form I laid out here) extend to logical grounding; and it can be indeterminate what metaphysically explains what. But to my mind these costs are acceptable, and I would dispute that they are even costs at all.

For many years, metaphysicians working on grounding and explanation proceeded in near total isolation from the extensive and sophisticated body of work on scientific explanation. This has begun to change recently as the field is entering its mature phase (and as also evidenced by my overview of the recent literature in Sect. 1). However, the prospects of Metaphysical Unificationism have not yet been assessed systematically. In this paper I have argued that we should welcome it as an attractive, powerful and hitherto unduly neglected newcomer to the growing number of competing theories of metaphysical explanation.

References

Achinstein, P. (1983). *The nature of explanation*. New York: Oxford University Press.

Armstrong, D. M. (1978). *Universals and scientific realism: A theory of universals*. Cambridge: Cambridge University Press.

Audi, P. (2012). A clarification and defense of the notion of grounding. In F. Correia & B. Schnieder (Eds.), *Metaphysical grounding* (pp. 101–121). Cambridge: Cambridge University Press.

Baysan, U. (2018). Epiphenomenal properties. *Australasian Journal of Philosophy, 96,* 419–431.

Bennett, K. (2017). *Making things up*. Oxford: Oxford University Press.

Berker, S. (2018). The unity of grounding. *Mind, 127,* 729–777.

Bertrand, M. (2018). Metaphysical explanation by constraint. *Erkenntnis.* https://doi.org/10.1007/s10670-018-0009-5

Correia, F. (2010). Grounding and truth-functions. *Logique et Analyse, 53,* 251–279.

Correia, F. (2013). Metaphysical grounds and essence. In M. Hoeltje, B. Schnieder, & A. Steinberg (Eds.), *Varieties of dependence*. München: Philosophia Verlag.

Correia, F. (2014). Logical grounds. *Review of Symbolic Logic, 7,* 31–39.

Correia, F., Schnieder, B., & (Eds.). (2012). Grounding: an opinionated introduction. In *Metaphysical grounding: Understanding the structure of reality* (pp. 1–36). Cambridge: Cambridge University Press.

Cowling, S. (2013). The modal view of essence. *Canadian Journal of Philosophy, 43,* 248–266.

Crane, T. (2001). *Elements of mind*. Oxford: Oxford University Press.

Daly, C. (2012). Scepticism about grounding. In F. Correia & B. Schnieder (Eds.), *Metaphysical grounding: Understanding the structure of reality* (pp. 81–100). Cambridge: Cambridge University Press.

Dasgupta, S. (2014a). On the plurality of grounds. *Philosophers' Imprint, 14*(20), 1–28.

Dasgupta, S. (2014b). The possibility of physicalism. *Journal of Philosophy, 111,* 557–592.

Dasgupta, S. (2017). Constitutive explanation. *Philosophical Issues, 27,* 74–97.

deRosset, L. (2013). Grounding explanations. *Philosophers' Imprint, 13*(7), 1–26.

deRosset, L. (2019). *Fundamental things*. Oxford: Oxford University Press (under contract).

Dorr, C., & Hawthorne, J. (2013). Naturalness. *Oxford Studies in Metaphysics, 8,* 3–77.

Dunaway, B. (ms). Perfectly natural relative naturalness. http://www-personal.umich.edu/~dunaway/RelativeNaturalness.pdf. (unpublished manuscript).

Fine, K. (1994). Essence and modality. *Philosophical Perspectives, 8,* 1–16.

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

Fine, K. (2010). Some puzzles of ground. *Notre Dame Journal of Formal Logic, 51,* 97–118.

Fine, K. (2012). Guide to ground. In F. Correia & B. Schnieder (Eds.), *Metaphysical grounding: Understanding the structure of reality* (pp. 37–80). Cambridge: Cambridge University Press.

Fine, K. (2015). Unified Foundations for Essence and Ground. *Journal of the American Philosophical Association,* 1, 296–311.

Friedman, M. (1974). Explanation and scientific understanding. *Journal of Philosophy, 71,* 5–19.

Hempel, C. G., & Oppenheim, P. (1948). Studies in the logic of explanation. *Philosophy of Science, 15,* 135–175.

Hirsch, E. (2008). Language, ontology, and structure. *Nous, 42,* 509–528.

Jenkins, C. S. (2013). Explanation and fundamentality. In M. Hoeltje, B. Schnieder, & A. Steinberg (Eds.), *Varieties of dependence: ontological dependence, grounding, supervenience, response-dependence* (pp. 211–241). Munich: Philosophia Verlag.

Jones, T. (1995a). Reductionism and the unification theory of explanation. *Philosophy of Science, 62,* 21–30.

Jones, T. (1995b). How the unification theory of explanation escapes asymmetry problems. *Erkenntnis, 43,* 229–240.

Kim, J. (1994). Explanatory knowledge and metaphysical dependence. *Philosophical Issues, 5,* 51–69.

Kim, J. (1998). *Mind in a physical world: An essay on the mind-body problem and mental causation.* Cambridge: MIT Press.

Kitcher, P. (1981). Explanatory unification. *Philosophy of Science, 48,* 507–531.

Kitcher, P. (1989). Explanatory unification and the causal structure of the world. In P. Kitcher & W. Salmon (Eds.), *Scientific explanation.* Minneapolis: University of Minnesota Press.

Kment, B. (2006). Counterfactuals and the analysis of necessity. *Philosophical Perspectives,* 20, 237–302.

Koslicki, K. (2015). The coarse-grainedness of grounding. *Oxford Studies in Metaphysics, 9,* 306–344.

Kovacs, D. M. (2017). Grounding and the argument from explanatoriness. *Philosophical Studies, 174,* 2927–2952.

Kovacs, D. M. (2018). What is wrong with self-grounding? *Erkenntnis, 83,* 1157–1180.

Lange, M. (2014). Aspects of mathematical explanation: Symmetry, unity, and salience. *Philosophical Review, 123,* 485–531.

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

Lewis, D. (1986a). *The plurality of worlds.* Oxford: Blackwell.

Lewis, D. (1986b). Causal explanations. In *Philosophical papers* (Vol. 2). Oxford: Oxford University Press.

Lipton, P. (2001). What good is an explanation? In G. Hon & S. S. Rakover (Eds.), *Explanation: Theoretical approaches and application.* Dordrecht: Springer.

Litland, J. (2017). Grounding ground. *Oxford Studies in Metaphysics, 10,* 279–316.

Marshall, D. (2015). Humean laws and explanation. *Philosophical Studies, 172,* 3145–3165.

Maurin, A.-S. (2018). Grounding and metaphysical explanation: It's complicate. *Philosophical Studies.* https://doi.org/10.1007/s11098-018-1080-0

Nagel, E. (1961). *The structure of science.* Harcourt: Brace & World.

Raven, M. (2013). Is ground a strict partial order? *American Philosophical Quarterly, 50,* 191–199.

Raven, M. (2015). "Ground", philosophy. *Compass, 10,* 322–333.

Raven, M. (2016). Fundamentality without foundations. *Philosophy and Phenomenological Research, 93,* 607–626.

Rosen, G. (2010). Metaphysical dependence: Grounding and reduction. In B. Hale & A. Hoffman (Eds.), *Modality: Metaphysics, logic, and epistemology.* Oxford: Oxford University Press.

Roski, S. (2017). Bolzano and Kim on grounding and unification. *Synthese.* https://doi.org/10.1007/s11229-017-1593-7

Ruben, D.-H. (1990). *Explaining explanation.* London: Routledge.

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

Schaffer, J. (2010). Monism: The priority of the whole. *Philosophical Review, 119,* 31–76.

Schaffer, J. (2016). Grounding in the image of causation. *Philosophical Studies, 173,* 49–100.

Schaffer, J. (2009). On what grounds what. In D. Chalmers, D. Manley, & R. Wasserman (Eds.), *Metametaphysics: New essays on the foundations of ontology* (pp. 347–383). Oxford: Clarendon Press.

Schnieder, B. (2010). A puzzle about 'Because'. *Logique et Analyse, 53,* 317–343.

Sider, T. (2003). Reductive theories of modality. In M. J. Loux & D. W. Zimmerman (Eds.), *Oxford handbook of metaphysics*. Oxford: Oxford University Press.

Sider, T. (2011). *Writing the book of the world*. Oxford: Oxford University Press.

Sider, T. (1993). *Naturalness, Intrinsicality, and Duplication*. Ph.D. Dissertation, UMass Amherst.

Skiles, A. (2015). Against grounding necessitarianism. *Erkenntnis, 80,* 717–751.

Smart, J. J. C. (1959). Sensations and brain processes. *Philosophical Review, 68,* 141–156.

Strevens, M. (2008). *Depth: An account of scientific explanation*. Cambridge, MA: Harvard University Press.

Taylor, E. (2018). Against explanatory realism. *Philosophical Studies, 175,* 197–219.

Thompson, N. (2016). Grounding and metaphysical explanation. *Proceedings of the Aristotelian Society, 116,* 395–402.

Trogdon, K. (2013). Grounding: Necessary or contingent? *Pacific Philosophical Quarterly, 94,* 465–485.

van Fraassen, B. C. (1980). *The scientific image*. Oxford: Clarendon Press.

van Gulick, R. (2001). Reduction, emergence, and other recent options on the mind/body problem: A philosophic overview. *Journal of Consciousness Studies, 8,* 1–34.

van Riel, R. (2013). Identity, asymmetry, and the relevance of meanings for models of reduction. *British Journal for the Philosophy of Science, 64,* 747–761.

Williams, J. R. G. (2007). The possibility of onion worlds. *Australasian Journal of Philosophy, 85,* 193–203.

Wilsch, T. (2015). The nomological account of ground. *Philosophical Studies, 172,* 3293–3312.

Wilsch, T. (2016). The deductive-nomological account of metaphysical explanation. *Australasian Journal of Philosophy, 94,* 1–23.

Wilson, J. M. (2014). No work for a theory of grounding. *Inquiry, 57,* 1–45.

Woodward, J. (2003). *Making things happen: A theory of causal explanation*. Oxford: Oxford University Press.