

The Metaphysics of Moral Explanations¹

Daniel Fogal (NYU) and Olle Risberg (Uppsala)

Draft as of January 2019 – Comments and feedback welcome!

Abstract. It's commonly held that particular moral facts are explained by 'natural' or 'descriptive' facts, though there's disagreement over how such explanations work. We defend the view that general moral principles also play a role in explaining particular moral facts. More specifically, we argue that this view best makes sense of some intuitive data points, including the supervenience of the moral upon the natural. We consider two alternative accounts of the nature and structure of moral principles—the nomic view' and 'moral platonism'—before considering in what sense such principles obtain of necessity.

1 Introduction

What are the aims of first-order moral inquiry, or normative inquiry more generally? One aim is to specify which actions are right and which actions are wrong. Arguably, though, that's not enough; another aim is to *explain why* the right actions are right and the wrong ones are wrong. But how, more exactly, do such explanations work?

On one natural view, the full explanation of why a particular action was wrong (or right or good or bad, etc.) involves two kinds of facts: (i) a particular 'natural' or 'descriptive' fact—perhaps it was a lie, for instance, or failed to treat humanity as an end in itself—and (ii) a general moral fact—that it is wrong to lie, for instance, or to fail to treat humanity as an end in itself. Call the latter facts *moral principles*.

The view that moral principles are explanatory in this way has several virtues. It makes sense of the search for general, first-order explanatory principles that many ethicists engage in. It also accords with plausible theories about other kinds of explanation, such as causal ones, which also seem to involve a general or 'law-like' element. And it accounts for the strangeness of statements like: 'Although Joyce acted wrongly because she lied, there's nothing wrong with lying itself.'

Despite these attractive features, however, the precise nature and structure of the relevant principles, and the form of explanation involved, is

¹ Forthcoming pending final peer review in *Oxford Studies in Metaethics*, Vol. 15. Thanks to Selim Berker, Matti Eklund, David Faraci, Martin Glazier, Jaakko Hirvelä, John Keller, David Mark Kovacs, Henrik Rydén, Jonathan Shaheen, Knut Skarsaune, Bart Streumer, Mark van Roojen, Pekka Väyrynen, Daniel Wodak, and the audiences at Chapel Hill, College of William & Mary, and University of Wisconsin-Madison for helpful feedback on earlier drafts of this paper.

not yet well understood. What's more, the view that moral principles play an explanatory role has recently been attacked. Selim Berker (2018b), for example, argues that moral principles can instead be viewed as *mere summaries* of the explanatory relations that obtain between particular moral and non-moral facts, with the principles themselves being explanatorily inert. Similarly, Mark Schroeder (2005; 2007) argues that Ralph Cudworth's objection against theological voluntarism—which crucially involves appeal to explanatory moral principles—threatens to generalize so as to rule out the possibility of “perfectly general explanatory moral theories” (2005, 3).

Our aim in this paper is to defend the explanatory role of moral principles by arguing that it best accommodates some intuitive claims about moral metaphysics. Indeed, while we focus on morality, parallel claims are plausible with respect to normative principles and explanations more generally. We begin by presenting the relevant ‘data points’ (§2) before outlining, in general terms, how our favored view captures them (§3). We then argue that the view that moral principles aren't explanatory in this way falls short (§4)—in particular, it fails to explain the supervenience of moral facts on natural facts. We proceed to discussing two competing accounts of what, more precisely, moral principles are like: *the nomic view* (§5) and *moral platonism* (§6). Finally, we consider in what sense moral principles are metaphysically necessary (§7).

2 The data

We'll begin by presenting three claims we think any fully satisfactory metaethical theory should explain, or otherwise accommodate in a principled manner. Although none are completely uncontroversial, theories that explain them will, other things being equal, enjoy an advantage over those that don't (and even more so over theories which are incompatible with them).

The first data point is that moral facts supervene upon purely natural (or purely descriptive or whatever) facts.² Suppose that Matti is a good person. Besides being good, he also has a host of natural properties (including relational ones) that are connected to goodness in the following way: anyone who is descriptively just like Matti will also be good. Indeed, if Matti is good, it's *impossible* for someone to possess all and only his natural properties without also being good. Generalizing, the supervenience relation that holds between ‘the natural’ and ‘the moral’ is standardly formulated as follows, where M is the family of moral properties, N is the family of natural (or whatever) properties, and \Box is metaphysical necessity:

² We say ‘purely descriptive’ since, at least assuming cognitivism, there's a clear sense in which moral judgments are also descriptive: they purport to describe moral facts. But insofar as they are also normative or evaluative, they are not *purely* descriptive.

Data point (i): Strong Supervenience

$$(\forall F \text{ in } M)(\forall x)[Fx \rightarrow (\exists G \text{ in } N)(Gx \ \& \ \Box(\forall y)(Gy \rightarrow Fy))]^3$$

In English: for every moral property F , if something is F , then that thing has some (possibly quite complex!) natural property G such that, by metaphysical necessity, everything that has G also has F .

Strong Supervenience has long been treated as something like a fixed point, though recently it's been called into question.⁴ We'll nonetheless assume that, defectors notwithstanding, its rejection comes as a cost.

There are at least three things to note about Strong Supervenience. The first is that despite being standardly glossed (as we did above) as expressing a relation between 'the moral' and 'the natural', the relevant pattern of covariation only concerns the moral and natural properties of *particular* things. (The higher-order quantifiers quantify over moral and natural properties while the first-order quantifiers quantify over particular bearers of those properties.) We'll return to this point below.

Second, there's controversy over how to characterize the relevant supervenience base. At a minimum, we assume the relevant family of properties (N) is closed under property conjunction and property disjunction and restricted to *repeatable* properties. Following Atiq (MS), a property B is repeatable when, necessarily, if x is B then it's metaphysically possible for there to be an entity y such that y is distinct from x and y is B . This rules out inclusion of haecceitistic properties, such as the property of being Barack Obama.⁵ More generally, the goal is to strip base properties of their particularity. There are a host of additional complications, however. Indeed, as Sturgeon (2009) notes, it's "difficult (...) to find a version of the doctrine of the supervenience of the evaluative that is available as a serious argumentative weapon in the dialectic of metaethics". Even so, we think

³ Cf. McPherson (2015), Leary (2017), Dreier (1992). It's standardly assumed that Strong Supervenience itself holds of necessity, either metaphysical necessity or conceptual necessity or both. We assume that it holds of metaphysical necessity; it may also hold of conceptual necessity, though we ourselves have too tenuous a grasp on that notion to take sides.

⁴ See, e.g., Rosen (forthcoming), Fine (2002), Hattiangadi (2018), and Roberts (2018). As Selim Berker (p.c.) points out, many deniers of Strong Supervenience remain committed to a form of supervenience—it's just one that involves 'normative' rather than metaphysical necessity. They thus face the analogous task of explaining why this supervenience claim holds. (For criticism of the notion of normative necessity, see Lange 2018) There are also extreme forms of particularism which are incompatible with supervenience. Arguably, however, such views are too extreme.

⁵ Atiq's restriction arguably isn't enough, however, since we also want to rule out repeatable properties like that of being an Obama, which Michelle and Barack share despite being numerically distinct. (Thanks to Selim Berker.)

everyone has to acknowledge that there is *some* truth in the neighborhood, even though how exactly one chooses to formulate it may depend on one's other commitments.

Third, Strong Supervenience doesn't entail that particular moral facts (e.g. that Matti is good) are *explained by* or obtain *in virtue of* natural facts. Indeed, supervenience claims in general are simply silent as to which, if any, explanatory relations obtain between the relevant kinds of facts.⁶ Nevertheless, the modal correlation specified by Strong Supervenience isn't the only interesting relation that holds between particular moral and natural facts. An explanatory relation also seems to hold: when a person is good or an action is wrong, for example, there are natural features of those entities that *make* the person good and the action wrong. The second data point concerns this relation:

Data point (ii): Particular moral facts are at least partly (and at least ultimately) explained by particular natural facts.

We say 'at least ultimately' because some particular moral facts may obtain partly in virtue of other particular moral facts. In such cases, we assume those other moral facts are in turn explained (again, at least ultimately) by particular natural facts.⁷

Some additional clarifications are in order. First, what is meant by 'particular moral fact'? A particular moral fact is a moral fact about a particular (i.e. dated, non-repeatable) thing, such as a particular action, person, or state of affairs. Besides particular moral facts, there are also general, non-particular moral facts, such as that pain is bad or that lying is pro tanto wrong. This intuitive distinction is easy to overlook. For example, Pekka Väyrynen (2013) writes in a related context:

[It] is very common to think that actions and other things have their normative and evaluative properties *in virtue of* their non-normative, non-evaluative properties. It is similarly very common for those who

⁶ On the difference between supervenience and explanation, see DePaul (1987), Bliss and Trogdon (2014), and Berker (2018a). For dissent, see Kovacs (forthcoming).

⁷ Some particular moral facts may resist such explanation, however. For example, consider the fact that Matti is such that if he tells a lie, then he does something *prima facie* wrong. Insofar as the explanation of such moral 'Cambridge facts' differs from the typical case, though, they seem to be exceptions to a general norm. Another particular moral fact that may not seem to be explained by natural facts is the fact that the fact lying violates the Categorical Imperative is a reason against Matti's lying yesterday. However, reasons-claims arguably require separate treatment, since they themselves are explanatory claims of a certain kind (cf. Fogal 2016), not entirely unlike claims about a particular action being pro tanto right or wrong in virtue of a given natural fact. (Thanks to Selim Berker for (versions of) these apparent counterexamples.)

are allergic to talk of normative properties nonetheless to agree that things are good or bad, or right or wrong, *because* of some non-normative properties. There is, in other words, a strong intuition that normative facts are dependent on and explained by other facts. Call this ‘the dependence intuition’. (p. 155; italics in original)

Though we agree with the spirit of the dependence intuition, this is a potentially misleading formulation of it insofar as it suggests that *every* moral (or normative) fact is explained by other facts. This matters because while *particular* moral facts plausibly depend (at least in part) on natural facts in this way, it’s far less clear that *general* moral facts—e.g., fundamental moral principles—also so depend. So while Väyrynen is right that there’s a strong intuition in the vicinity, it only applies to a subset of the moral facts—i.e., the particular ones. We’ll return to this issue below.

Second, what is meant by ‘explained’? The kind of explanation we have in mind is *metaphysical* explanation. One distinguishing feature of such explanations is that they are non-causal. Suppose, for example, that we want to know why the barn is red.⁸ There are at least two questions we might be interested in. The first is what *made it the case* that the barn is red. The fact that, say, someone painted it yesterday would help provide an answer—it would help causally explain why the barn is red. The second question, in contrast, concerns what presently *makes it the case* that the barn is red. Here historical facts are irrelevant. Instead, what matters are contemporaneous facts, such as the fact that the barn is crimson. This non-causally explains why the barn is red.⁹

In addition to being non-causal, we take metaphysical explanation to be *objective*, roughly in the sense of being mind- or stance-independent.¹⁰ Impressionistically put, objective explanations solely involve explanatory relations—whether they be causal, nomic, metaphysical, or something else—obtaining ‘out there’ in the world. They are to be contrasted with the more familiar, pragmatic notion of an explanation, understood as the sort of thing we standardly ask for and provide concerning a variety of subject matters (and which can in that sense be causal or non-causal), and whose success depends on “facts about the interests, beliefs or other features of the psychology of those providing or receiving the explanation [or] the ‘context’ in which the explanation occur” (Woodward 2014, §6.1). In slogan form, it’s important to distinguish between ‘explains’ in the sense of ‘makes it the case’

⁸ Cf. Glazier (2016, 11).

⁹ We count both present facts about the past and timeless facts as contemporaneous.

¹⁰ We treat this as a stipulation. It’s therefore immune to challenges by Thompson (2016) and Miller and Norton (forthcoming), who seem to collapse the ‘makes sense of why’ vs. ‘makes it the case’ distinction.

and ‘explains’ in the sense of ‘makes sense of why’.¹¹ Although knowing what made or makes something the case will usually help make sense of why that thing is the case, objective and pragmatic explanations are not neatly aligned.

It might be objected that focusing on metaphysical explanation prejudices certain metaethical disputes, such as whether some form of expressivism is true. Insofar as there’s a conflict, however, we think the problem is likely to lie with those views. After all, everyone needs a story about the metaphysical-seeming judgments we make about moral matters, including judgments about what makes actions right, wrong, and so on. If expressivists can provide such a story, then the seeming conflict disappears.¹² If they can’t, that’s a problem for their view.

The third clarification concerns what is meant by ‘partly explains’. The relevant contrast here is the notion of a *full* explanation. Almost all the explanations we actually provide—i.e. pragmatic explanations—are partial rather than full. If you purchase a pet, for instance, and someone asks you why, you might reply that you’re lonely. While sensible as a reply, the fact that you’re lonely doesn’t *by itself* explain why you purchased a pet. Instead, its explanatory potential depends on a bunch of background facts that are taken for granted, such as that you don’t want to be lonely, that you believe a pet will make you less lonely, and so on. Your loneliness is thus only a part, and in fact a rather small part, of what fully explains your action. This illustrates a general feature of our explanatory practice: rather than citing everything required to fully explain something, we’re typically content to highlight one or two particularly notable factors, trusting our interlocutors to fill in the rest. Providing the full explanatory story is typically laborious, if possible at all, and unnecessary.

Here, finally, is our third data point:

Data point (iii): Moral principles are explanatory.

This claim is motivated, in part, by the observation that a moral theory merely consisting of a list of all particular (actual and/or possible) actions that are right or wrong is incomplete. We also want to know *why* they made it on the list. General moral principles provide answers to that question.

Some care is needed here, however. For as Berker (2018b) notes, moral principles might be explanatory in either of two ways. Suppose, for example, it’s always wrong to lie. According to Berker, this principle is explanatory only in the sense that it specifies a natural property—in this case, being a lie—such that any particular act with that property is wrong fully because it is a lie. On

¹¹ More nuance is clearly called for here, though we lack to space to provide it. Suffice it to say, ordinary explanation-talk is both messy and context-sensitive (cf. Lewis 1986; Jenkins 2008).

¹² For discussion, see Berker (forthcoming) and Toppinen (2018).

this view, moral principles can be viewed as mere *summaries* of patterns of particular explanatory relations, with the principles themselves being explanatorily inert. As we'll put it, moral principles thus understood are merely *explanatory in content*—they specify which particular natural facts explain which particular moral facts, and that's it. The salient alternative, which we prefer, is that the truth of the principle that lying is wrong *itself* partly explains particular moral facts involving lying. The fact that the action was a lie is therefore only part of what fully explains why it's wrong. On this view, moral principles are *explanatory in role*—they themselves figure in the explanation of particular moral facts.¹³ Data point (iii) is neutral with respect to whether such principles are explanatory in role or in content.

3 Explaining the data

To recap, here are the three data points we'll focus on:

Data point (i): Strong Supervenience—for every moral property F, if something is F, then that thing has some natural property G such that, necessarily, everything that has G also has F.

Data point (ii): Particular moral facts are at least partly explained by natural facts.

Data point (iii): Moral principles are explanatory.

A unified account of these data points is desirable. The most obvious way of giving one is to view some as more fundamental than others. But which ones?

We doubt (i) is the most fundamental. For as Jaegwon Kim and others have emphasized, supervenience is “not a ‘deep’ metaphysical relation” but “a ‘surface’ relation that reports a pattern of property covariation, suggesting the presence of an interesting dependency relation that might explain it” (Kim 1993: 167). Supervenience theses thus *call for* explanation rather than provide them. So ideally (i) can be accounted for in terms of (ii) or (iii) or both.

We think the best way of making sense of (i) is by opting for the package of (ii) together with the interpretation of (iii) on which moral principles are explanatory in role. On this view, explanations of particular moral facts involve the following main ingredients:

Explanans: particular natural fact(s) (e.g. *a* is a lie)

Principle: general explanatory moral principle (e.g. Lying is wrong)

Explanandum: particular moral fact (e.g. *a* is wrong)

¹³ Berker calls such principles ‘explanation-serving’ and principles that are explanatory in content ‘explanation-involving’.

The explanandum is fully explained by the explanans together with the moral principle, though (as we'll see) the exact role played by principles will depend, among other things, on how they are formulated. But for now this outline of the tripartite structure of moral explanations will suffice.

It's worth noting that alternative terminological choices might be made. Schaffer (2017a), for example, emphasizes the importance of the tripartite structure of metaphysical explanations by distinguishing between the 'source', 'link', and 'result' in an explanation. As he notes, when "there are three roles involved, nothing but confusion can arise from insisting on only using two classificatory boxes" (20). Causal explanations also illustrate this point: while causal laws plausibly help explain particular effects, laws aren't themselves causes. Causes and laws play different roles in the full explanation of whatever is caused. We agree, although we've chosen to supplement the explanans/explanandum ideology rather than jettison it.

The tripartite account of moral explanations straightforwardly accommodates data points (ii) and (iii): particular moral facts are explained by particular natural facts together with general moral principles.¹⁴ Given the substantive but plausible assumption that the fundamental moral principles obtain of metaphysical necessity (if at all), making sense of strong supervenience—data point (i)—is also straightforward.¹⁵ We'll call this the *Divide & Conquer* (or *D&C*) strategy.

Strong Supervenience essentially states that the moral properties of some particular thing couldn't be different unless its natural properties were different. On the tripartite account, the moral properties of particular things depend on two things: (a) which natural properties they have and (b) which moral principles are true. Regarding (a), it's trivial that there can't be a

¹⁴ In this respect, our account resembles what Schroeder (2005; 2007; 2014) calls "the Standard Model" of normative explanations, where such explanations "[subsume] specific obligations in context to more general obligations, by appeal to specific features of the agent's circumstances" (2014: 3). Although we agree this is an intuitive idea, we think the tripartite account is more perspicuous than Schroeder's Standard Model. Especially important here is the distinction between particular and general normative facts (cf. sect. 2), since some of Schroeder's discussion seems to overlook it. For example, Schroeder apparently holds that every "perfectly general explanatory moral theory" is in the business of explaining why particular agents ought to perform certain action-types (2005, 3). What we perform in the first instance, however, are *token* actions, and the moral status of such actions is also—arguably even primarily—something moral theories should explain. Other times, something like the general-particular distinction figures in Schroeder's discussion of the Standard Model, but is not clearly separated from other distinctions, such as the one between explained and unexplained moral facts, or the one between moral facts that are 'wholly distinct' from non-moral facts and those that aren't (cf. 'pure' moral facts; 2014, ch. 6). Contra Schroeder, then, we doubt the Standard Model—at least as he formulates it—is widely accepted among moral philosophers. Reinterpreted along our lines, though, we think it's both widely accepted and highly plausible.

¹⁵ This assumption can be understood in different ways. We'll return to it in sect. 7.

difference in something's natural properties without there being a difference in its natural properties. And regarding (b), it's only slightly less trivial that there can't be a difference in which moral principles are true without there being a difference in a thing's natural properties—because there can't be a difference in the moral principles, period. They obtain of necessity and so trivially supervene on everything. So it's no surprise that there can't be a difference in a particular thing's moral properties without there being a difference in its natural properties.

We take the claim just explained to be the 'core' supervenience claim. But as David Faraci (2017) emphasizes, Strong Supervenience also entails a *necessitation* claim—namely, that no particular thing can have a moral property unless it also has some natural property. Making sense of this additional claim isn't particularly problematic. For its falsity would require it to be possible for something to have a moral property while lacking any and all natural or descriptive properties. Given the very broad sense of 'natural' and 'descriptive' at issue, however, it's hard to see what such a thing could be like. Every action, for example, will at least have the natural property of being an action, no matter how featureless it might otherwise be. Similarly, Cartesian souls, even if empirically inaccessible, will still have properties such as being conscious, or, indeed, being Cartesian souls. This suggests that every particular thing will have at least some descriptive property or other, and *a fortiori*, that every particular thing that has a moral property will have some descriptive property or other.¹⁶

At this point it might be worried that an account of Strong Supervenience that posits non-contingent explanatory moral principles fails to make genuine progress, since such principles merely repeat (in a slightly different guise) the problematic truth that was supposed to be accounted for.¹⁷ But that's false: supervenience theses say nothing about explanation, and so can't be repeated by claims about such.

Although several recent explanations of supervenience resemble the D&C strategy, none are fully satisfactory. Enoch (2011), for instance, takes the basic normative facts to be "norms" that hold with metaphysical necessity, and thereby explain Strong Supervenience. But Enoch says little about what norms are or how they explain particular normative facts. We'll consider some different answers to this question in sections 5-6.¹⁸ One of those accounts

¹⁶ We suspect that Faraci's real concern is why "it is impossible for normative properties to be ungrounded [or at least not grounded in natural properties]" (2017, 315). But this a concern about data point (ii) rather than data point (i). As we've emphasized, the supervenience thesis says nothing about what explains (or "grounds") what.

¹⁷ Cf. McPherson (2012) on "bruteness revenge" and McPherson (2015), sect 4.3.

¹⁸ Enoch returns to the issue in his (forthcoming), advocating a form of so-called 'grounding pluralism'. For a potentially—and we think actually—devastating objection,

takes the fundamental moral facts to concern *kinds* rather than particulars. Both Skarsaune (2015) and, more briefly, Schroeder (2014, sect. 6.6) appeal to this idea in explanations of supervenience, though neither pays sufficient attention to data points (ii) and (iii). Scanlon (2014), in contrast, holds that supervenience is explained by “pure” normative facts of the form ‘ $\forall x(R(p, x, c, a))$ ’, which reads: for all agents x , in circumstance c , fact p is a reason to do a . However, for reasons discussed in section 4, the use of universal generalisations in formulating moral principles is problematic.¹⁹

While the proposals just considered are all pursued in defense of some form of ‘non-naturalism’ about moral facts, the D&C strategy itself is neutral with regard to the naturalism/non-naturalism dispute, and indeed with regard to most other disputes in metaethics. In particular, even if it’s always possible to explain why a particular thing has a certain moral property, it doesn’t follow that the moral properties themselves—that is, *what it is* to have a certain moral property—can be explained in any way.²⁰ That’s an important point of contention dividing naturalists and non-naturalists, and one to be decided on other grounds.²¹ More generally, the D&C strategy is neutral concerning the status, specific content, number, complexity, and explanation (if any) of the relevant moral principles. We view this as a feature, not a bug: it focuses the debate precisely where it should be (and to some extent always has been)—namely, on the fundamental moral principles, if such there are.

Given what we’ve said so far, the attention garnered by Strong Supervenience (and its ilk) over the past few decades begins to look misplaced. It has frequently been held, for example, that non-naturalists have an especially hard time explaining moral supervenience. But if the D&C strategy is successful, this is a mistake. This is just one example of how focusing on supervenience while neglecting the more fundamental data points (ii) and (iii) can be distorting. Other examples are furnished by various attempted explanations of supervenience—many of which are non-naturalist—that are either silent about or, worse, in tension with the other data points. To illustrate:

see Enoch (forthcoming), n. 21. In our view, the best way to develop Enoch’s 2011 account is instead along the lines that we’ll consider in section 5.

¹⁹ Schroeder (2015) suggests that the quantifier and the person variable x are redundant, and that Scanlon’s proposal can be read ‘ $R(p, c, a)$ ’ without loss. And indeed, that is the formulation Scanlon uses in the manuscript for his 2009 Locke Lectures, on which the book is based. Thus understood, Scanlon’s view is a version of platonism, which we discuss in section 6.

²⁰ As Leary (ms) argues, this distinction was recognized by G.E. Moore (1942, 588). See also Rydén (2018) on “opaque grounding” (which we’ll return to below) and Rosen (2010: sect. 13) on “Moorean connections”.

²¹ Thus, the distinction between explaining what it is to be F and explaining why something is F helps avoid Berker’s revisionary result that “almost all contemporary metaethical views (other than nihilism) end up counting as a form of non-naturalism about the normative” (2018b: 29) and Heathwood’s (2012) related view.

- Naturalistic identity-theses make sense of supervenience but are harder to reconcile with the data points about moral explanation. Of course things that differ morally must differ naturally if moral properties just *are* natural properties. But if rightness is, say, identical to happiness-maximizing, it's hard to see how actions can be right *in virtue of* being happiness-maximizing, since nothing explains itself (cf. McNaughton & Rawling 2003).
- Expressivist explanations of supervenience have often centered on the idea that our moral practice wouldn't fulfill its 'function' of, e.g., coordinating behavior, if our moral views didn't respect supervenience. For example, Mitchell (2017) attributes to Blackburn the view that "it is practically necessary that everyone conform their evaluations to a supervenience constraint: without it, we lapse into practical and discursive chaos...". Whether or not this thesis makes sense of supervenience, however, it's silent with respect to moral explanations, and hence data points (ii) and (iii).

There are various things adherents of such views might say in response.²² But our main complaint concerns the focus of the debate. For as Berker (2018a) convincingly argues, the notion of moral supervenience was first introduced and motivated with reference to claims concerning *both* co-variance—data point (i)—and explanation—data point (ii). In one of the earliest discussions, for example, R. M. Hare notes that if two things differ in goodness "there must be some further difference between them to *make* one good and the other not..." (1952, 81; emphasis added). The modal formalizations of supervenience that subsequently became dominant replaced this dual-focus with a single-minded one—they express claims about co-variation and that's it. This coincided with a general philosophical suspicion of heavier-weight notions like metaphysical explanation, but times have changed—and for the better, to our minds. Both of Hare's original motivations remain relevant, and accounts that don't make sense of moral explanations fail to capture what motivated the focus on supervenience in the first place.

²² For example, in order to make progress in explaining data point (ii), expressivists might help themselves to the account of normative explanation suggested by Berker (forthcoming)—of which both he and we are ultimately skeptical. It's also possible to simply deny one of our data points and seek to explain away the relevant intuition(s). As noted above, however, we think such views will be less plausible, other things being equal. For example, while our concern has been the broadly metaphysical sense of 'explains', there may be other senses of the word that advocates of identity theses can invoke if they want to say that some facts explain themselves. For instance, one might say 'Joe is a bachelor because he is an unmarried eligible male' in attempting to elucidate the notion of a bachelor (see Kovacs 2018, sect. 4). The present point is just that moral explanations don't strike us as mere attempts at elucidation. We'll return to the relation between explanatory and identificational claims in section 6.1.

To sum up so far: taking moral principles to be explanatory in role allows us to provide a unified account of all three data points. The next task is to provide an account of the nature and structure of such principles. Doing so is not straightforward—for as we’ll see, the most common way to formulate moral principles makes it hard to see how they could be explanatory in role, rather than merely in content.

4 Moral principles are not merely explanatory in content

Moral principles are commonly formulated as universal generalizations. Berker (2018b) defends this view, arguing that “the most naive way of formulating moral principles (...) is also the best, namely one that uses no materials other than a wide-scope necessity operator, standard quantification, mundane indicative conditionals, and the full grounding relation” (26).²³ The utilitarian principle is thus formulated as follows (‘B’ for ‘Berker’):

U_B: Necessarily, an action is required if and only if, and fully because, it maximizes happiness. (2018b, 26)

It will be useful to pay attention to some details of this formulation. While Berker frequently says statements like this one are necessitated universal generalizations, or necessitated universally generalized indicative (bi)conditionals, U_B doesn’t explicitly involve universal generalization as it stands. Rather, it uses the indefinite ‘an action’ which (as we’ll further discuss in section 5.1) can be used in many different ways. Thus, for now, we’ll assume that U_B can be paraphrased as follows:

U_B*: Necessarily, for all actions *x*, *x* is required if and only if, and fully because, *x* maximizes happiness.

U_B* states a general fact—call it ‘[U_B*]’—that is explanatory in content, but not in role: it states that whenever some particular action is morally required, that fact is fully ‘grounded’ in its maximizing happiness.²⁴ Given the common assumption that universal generalizations obtain at least partly in virtue of their instances, [U_B*] obtains (if at all) in virtue of *this* action being required because it maximizes happiness, *that* action being required it maximizes happiness, and so on for each possible required action (perhaps together with

²³ Berker assumes that so-called metaphysical grounding is “itself an explanation relation” (2018b, 4), so that statements of what metaphysically explains what can also be made in terms of ground.

²⁴ Following Berker (following Rosen), we’ll use square brackets to denote facts. We’ll also sometimes use them to denote fact schemata—context will disambiguate.

a totality fact).²⁵ $[U_B^*]$ thus summarizes, and is explained by, instantiations of the grounding relation between particular facts about happiness-maximizing and particular facts about moral obligatoriness— $[U_B^*]$ itself plays no explanatory role.²⁶ Indeed, given that metaphysical explanations are irreflexive and transitive, $[U_B^*]$ can't explain the relevant particular facts about moral obligatoriness.

The view suggested by Berker thus resembles 'Humean' accounts of laws of nature, according to which natural laws are mere summaries of patterns among particulars. Yet the view is also *non*-Humean given that the relevant patterns involve instantiations of a hyperintensional explanatory relation, rather than mere co-occurrence of distinct properties or facts. This view—call it *Hyperintensional Humeanism*—is effectively what you get if you take data point (ii) concerning the explanation of particular moral facts by natural facts to be more fundamental than data point (iii) concerning the explanatory nature of moral principles.

The problem, however, is that Hyperintensional Humeanism has trouble making sense of the supervenience of the moral upon the natural—i.e. data point (i). To see why, let D be the set of Matti's natural properties. Assuming he's good, Strong Supervenience allows us to infer that every other possible entity with the properties in D is also good. And if we take (ii) to be more fundamental than (iii), it's not clear why that inference is sound. On Berker's view, there are some natural properties in D such that the fact that Matti has those properties fully explains the fact that he's good. Call those natural properties D^* . Given so-called *grounding necessitarianism*, it follows that, necessarily, if Matti is D^* then he is good.²⁷ But it doesn't follow that if, say, *Folke* has the properties in D (and thus the ones in D^*) then *Folke* is good. That's because Berker takes the grounding relation to hold between wholly particular natural facts (e.g., that Matti is D^*) and wholly particular moral facts

²⁵ It's actually a hard question what the 'instances' of U_B^* are. Rather than facts of the form $[x$ is required because a is happiness-maximizing], it might be better to view them as particular facts of the form $[x$ is required if and only if and because a is happiness-maximizing]. This would complicate things, especially if such facts can obtain even if the relevant action is neither happiness-maximizing nor required. But those complications have little bearing on the question of whether Berker's view makes sense of supervenience.

²⁶ Contra Rosen (2017), we think this is true even when the relevant universal generalizations are 'non-accidental'. (Rosen focuses on formulations of laws that have the form $\Box \forall x(Gx \rightarrow Fx)$, where G is a descriptive property, F is a normative property, the arrow stands for material conditional, and the box stands for so-called 'normative necessity'.)

²⁷ Grounding necessitarianism is the thesis that if some facts Γ fully explain the fact $[Q]$, then it's necessary that if the facts in Γ obtain then $[Q]$ obtains. This thesis is not universally accepted (see, e.g., Leuenberger 2014 and Skiles 2015). But in this context, rejecting necessitarianism would only make it more difficult to see how Berker's view could make sense of why Strong Supervenience is true.

(e.g., that Matti is good). As a result, nothing entitles us to generalize from facts about the natural and moral properties of one particular entity to facts about the natural and moral properties of another. In other words, even if the *properties* in *D* are repeatable, the *subjects* instantiating them—i.e, Matti and Folke—remain particular and non-repeatable. And generalizations from facts about one particular to another is precisely what that the supervenience thesis captures: if someone with the properties in *D* is good, then *anyone* with those properties is good as well.

Of course, *if* a principle akin to U_B^* is true for goodness, then it follows that goodness supervenes on the natural properties specified by that principle. But the point is that given Berker's underlying metaphysics, there's no reason to expect there to be true principles of that form, and hence no reason to expect supervenience to be true. It remains to be seen whether going from the wholly particular to the general will result in such principles—it all depends on what the Humean mosaic in all possible worlds turns out to look like. Thus, whether we get the Hyperintensional Humean's principles (and thus supervenience) is hostage to a kind of modal miracle.

In other words, although principles like U_B^* , if true, would secure supervenience, they would do so in the wrong way—rather than being made sense of in a principled fashion, it would still look like a mystery that the Humean mosaic necessarily turns out to give rise to such principles. While this worry resembles some of the traditional complaints in the literature on moral supervenience, the present situation is in one regard even worse.²⁸ For while the traditional worry is that it's mysterious why particular moral facts should necessarily align with particular natural facts in the way that Strong Supervenience specifies, the worry now concerns not only those facts but also particular instances of the grounding-relation between them.

In response, the Hyperintensional Humean might appeal to a principle like *Formality* (Rosen 2010, 131). Simplifying somewhat, this principle states that [*a* is *F*] fully grounds [*a* is *G*] only if any other fact of the form [*x* is *F*] fully grounds [*x* is *G*]. Given *Formality*, the fact that Matti is good because he is *D** entails that Folke is also good because he is *D**. However, this principle seems susceptible to the same problem. For why should the Hyperintensional Humean expect it to be true? Given her other commitments, she cannot consistently claim that this general principle (or something like it) is *itself* explanatory in role. Instead, whether *Formality* is true also depends on what the various patterns among wholly particular facts across possible worlds

²⁸ For example, see Skarsaune (2015, 267). There are also other worries about supervenience which we don't have in mind here—for instance, the one associated with Blackburn (1984) that it's mysterious why moral facts *weakly* supervene upon natural facts without *strongly* doing so. As Shoemaker (1987) rightly notes, this seems to be a “nonfact”.

happen to look like. Hence, rather than vindicating Strong Supervenience, Formality seems hostage to the same sort of modal miracle.

These problems are avoided if moral principles are explanatory in role. On this view, Matti's being D^* doesn't fully explain why he's good—a general principle connecting the property D^* to goodness also plays a role. If particular facts about goodness are always partly explained by such principles, it's clear why the assumption that Matti is both D and good allows us to infer that Folke, who is also D , must likewise be good. So we not only get Strong Supervenience (as already mentioned) but also a counterpart of Formality—at least in the moral case—for more or less the same reasons.²⁹

Of course, the arguments just given rely on several assumptions which could be rejected: for example, that (necessitated) universal generalizations are explained by their (non-necessitated) instances, that U_B is correctly paraphrased as a universal generalization in the first place, and so on. But once these assumptions are rejected, it becomes unclear what the Hyperintensional Humean view about moral principles is exactly, and whether the main ingredients in principles like U_B —i.e. “standard” quantification, indicative conditionals, and so on—are really as “naive” as Berker claims them to be. We won't consider those questions here. Instead, in what follows we'll focus on two accounts which allows moral principles to be explanatory in role, rather than merely in content.

5 The nomic view of moral explanations

The nomic view of moral explanations is modeled on what we'll call the law-based view of grounding explanations more generally. By 'grounding explanations' we mean cases in which a particular fact obtains in virtue of others (its 'grounds'), with the latter being more fundamental than the former. Standard examples include:

- (i) Mental facts obtain because of neurophysiological facts (Correria & Schneider 2012)
- (ii) The fact that the ball is red and round obtains in virtue of the fact that it is red and the fact that it is round (Fine 2012)
- (iii) Socrates was pale because he was this specific skin tone (Schaffer 2017b)

The nature of grounding explanations is a matter of controversy. One issue concerns the existence of general explanatory metaphysical principles, or 'laws', and what role (if any) they play. According to the law-based view,

²⁹ In support of Formality, Rosen notes that it seems “particular grounding facts must always be subsumable under general laws... It would be interesting to know why this is so.” (2010, 132) The idea that the general laws are explanatory in role is, we think, a natural answer to this question.

metaphysical laws play an ineliminable role in grounding explanations. On this view, the full metaphysical account of, say, Socrates' being pale extends beyond his having a certain specific skin tone—the general fact that having that skin tone makes one pale is also relevant.

Even if metaphysical laws exist and play an explanatory role, the precise role they play is subject to dispute. Should we treat laws as *part* of grounding explanations, for example, or instead as merely 'underlying' or 'backing' or 'governing' such explanations? (And what's the difference?) We'll set such questions aside because what matters for our purposes is what unites law-based theorists—namely, that metaphysical laws play some essential role in metaphysical explanations. The law-based view of grounding explanations can thus be seen as a generalized version of the tripartite view of moral explanations above:

Grounds: particular fact(s).

Metaphysical Law: general explanatory principle.

Explanandum: particular fact.

Discussions about grounding explanations are complicated, however, by the fact that the word 'ground(s)' is often used in different ways. For some, for A to ground B is just for A to metaphysically explain B,³⁰ whereas for others, for A to ground B is for there to be a relation between A and B that 'backs' metaphysical explanations.³¹ What's more, many theorists—including Berker—take grounding to be a relation holding solely between particular facts. But on the law-based view, it's a mistake to focus exclusively on relations between particular facts, just as on law-based views of causal explanation it's a mistake to focus exclusively on relations between individual events. In both cases laws also play a part. To minimize confusion we'll mostly avoid using 'ground(s)' as a verb, preferring instead to use it as a noun to pick out the explanan(s) of grounding explanations (per above) on analogy with the use of 'cause(s)' as a noun. This terminological issue will re-arise in responding to an objection by Berker in Section 5.1.

While the law-based view of grounding explanations is plausible, we won't be defending it here.³² Our primary concern is instead with the nomic view of moral explanations. What motivates the nomic view is the idea that the general principles or 'laws' that partially explain moral facts should be understood in the same way as those thought (rightly or wrongly) to figure in grounding explanations elsewhere. The main difference is that moral laws involve moral properties whereas non-moral laws do not—a difference in

³⁰ See, e.g., Litland (2013), Dasgupta (2014) and Wilsch (2015).

³¹ For example, see Audi (2012).

³² Instead, see, e.g., Kment (2014), Wilsch (2015, 2016), Glazier (2016), and Schaffer (2017a, 2017b).

content rather than form. Although the nomic view is thus a natural companion to the law-based view, neither entails the other. One might accept the nomic view, for example, while rejecting the law-based view in general; alternatively, one might accept the law-based view while rejecting the nomic view (by, e.g., denying that moral explanations are grounding explanations).

How, then, are metaphysical and/or moral laws to be understood? Although there are different ways of trying to capture their characteristic features, such as their generality and directionality, the proposal we find most congenial is from Glazier (2016):

[A metaphysical law] clearly has a sort of generality, but it is a general fact that is not explained by its instances. Since this sort of generality is not achieved through quantification, it must instead be achieved through another variable-binding operator. I therefore propose that we recognize a new operator ‘<<’ [that] bind[s] any number of variables, since our intuitive understanding of a general metaphysical-explanatory connection does not support any relevant limit. And because a fact may be metaphysically explained by any number of other facts, the operator should also be variably polyadic ‘on the left’. (25)

Thus, according to Glazier, a statement of a law will be of the form

$$\phi_1, \dots, \phi_n \ll_{\alpha_1, \dots, \alpha_m} \psi$$

where $\phi_1, \dots, \phi_n, \psi$ are sentences and $\alpha_1, \dots, \alpha_m$ are variables. We can therefore express the law connecting, say, an arbitrary thing’s being crimson and its being red as:

$$\text{Crimson:} \quad x \text{ is crimson} \ll_x x \text{ is red}$$

In terms of facts: *Crimson* states the general explanatory connection that holds between facts of the form [x is crimson] and facts of the form [x is red]—a connection that non-vacuously obtains even in possible worlds in which nothing is either crimson or red.

Though Glazier doesn’t discuss moral principles, the nomicist can use the ‘<<’-operator to formulate the utilitarian principle as follows (‘N’ for ‘Nomic’):

$$\text{U}_N: \quad x \text{ maximizes happiness} \ll_x x \text{ is morally required}$$

U_N states the general explanatory connection that holds between facts of the form [x maximizes happiness] and facts of the form [x is morally required]. Other principles can be formulated similarly. The nomic view is thus an

independently-motivated implementation of the tripartite view of moral explanations presented in section 3, thereby accommodating data points (i)-(iii) as desired.

5.1 Berker's objections

Berker (2018b) considers but rejects the nomic view of moral explanations. One worry concerns the relationship between metaphysics and natural language. He writes:

[It] is basically impossible to express [Glazier's] operator in natural language. Such a result would be surprising (...) Natural language is like water: over time it tends to adjust itself and flow toward the metaphysically and normatively relevant cracks in nature. (24)

Following this analogy, Berker urges "caution when there is such a large disconnect between a posited fundamental feature of reality and the everyday language with which we talk about that reality." (24)

We take the analogy with water to be inapt. Carnap (1963) provides a better analogy:

[N]atural language is like a crude, primitive pocket knife, very useful for a hundred different purposes. But for certain specific purposes, special tools are more efficient (...) If we find that the pocket knife is too crude for a given purpose and creates deficient products, we shall try to discover the cause for the failure, and then either use the knife more skillfully, or replace it for this special purpose by a more suitable tool, or even invent a new one. (938)

That is, natural language has many virtues, but for specialized purposes—like describing 'fundamental features of reality'—it often fails to provide the tools we need. Fortunately, its adaptability and flexibility enables us to give voice to novel concepts easily, whether by introducing a new term or else by giving a new meaning to an old term. In neither case does natural language adjust itself so as to "flow toward the metaphysically relevant cracks in nature"—*we* adjust it to do that.

The same is true of the generality that is characteristic of laws. While Berker may be right that "English doesn't mark a difference between the sort of generality featured in ['That an action breaks a promise makes it the case that it is *prima facie* wrong'] and the sort featured in a sentence such as 'That an action breaks a promise is not something its agent always knows'" (2018b: 25), we don't take this to be an obstacle or an objection. Indeed, it's a familiar point that simplicity of surface grammar often masks a multitude of ways in

which a sentence can be used. For example, sentences of the form ‘A(n) F is G’ can plausibly be used to express universal generalizations (‘A human cannot fly’), exception-permitting generalizations (‘A dog has four legs’), claims about specific individuals (‘A cat is on my lap’), and full or partial definitions (‘A bachelor is an unmarried male’, ‘A vixen is a fox’). The indefinite (‘an F’) is also used to express variable-binding term operators like Hilbert’s ε (cf. Woods 2014), and arbitrary reference more generally (e.g. ‘Let John be a German’; cf. Breckenridge & Magidor 2012). So: we can do lots of stuff with indefinites. Definitions are particularly relevant, since—as we’ll see in Section 6.1—they seem to involve the same kind of generality that laws do and similarly require special-purpose operators for perspicuous representation.

Berker has a second objection. As he notes, “Glazier often summarizes his proposed laws using ‘makes the case’-talk, which he takes to be a ‘generic placeholder’ that ‘perhaps... should be replaced by something like ‘determines’ or even ‘grounds’” (Berker 2018b, 24, citing Glazier 2016, 21, n. 15). If we take such glosses seriously, we might think that U_N can be reformulated as follows:

U_N^{**} : That an action maximizes happiness makes it the case that it is morally required.

However, if we take the ‘makes it the case’-locution in U_N^{**} to pick out the grounding-relation, Berker thinks that we face a dilemma. If this locution denotes *full* grounding, the moral principle becomes explanatorily redundant—the fact that the action maximized happiness by itself fully explains that it is right. On the other hand, if ‘makes it the case’ only denotes *partial* grounding, Berker thinks U_N^{**} ceases to qualify as a statement of moral law, since when looking for such laws “we are looking for [a] specification of *all* the partial grounds” (25). On this horn, Berker thinks, U_N is better informally put as follow:

U_N^{***} : That an action maximizes happiness, together with this very moral law, fully makes it the case that the action is morally required.

But this is metaphysically circular. And that’s implausible.³³ The upshot is that whether we take ‘makes it the case’ to denote full or partial grounding, the view that moral principles are explanatory in role runs into trouble.

There’s a lot that could be said in response to Berker’s dilemma. We’ll limit ourselves to three remarks.³⁴ To begin with, we need to be careful with ‘ground(s)’/‘makes-it-the-case’-talk. On the law-based view, grounding

³³ See Berker (2018b, sect. 4) for arguments to that effect.

³⁴ Thanks to Martin Glazier for discussion.

explanations have a tripartite structure: grounds, general explanatory principle, and explanandum. If ‘fully grounds’ means ‘constitutes the explanans and connection of’, then Berker is right that U_N doesn’t entail that A’s maximizing happiness fully grounds A’s being morally required. But U_N still legitimately ‘qualifies as a statement of moral law’ since it specifies all the relevant (i.e. particular-level) explanantia, or what we earlier called ‘grounds’. If, on the other hand, ‘fully grounds’ means ‘constitutes the explanans (but not the connection)’, then the entailment may well hold. But we’ll still need an explanatory connection, and we don’t see why U_N couldn’t state that connection without redundancy.

Regarding the second horn, we think the nomicist is within their rights taking the notion of metaphysical law as an ideological primitive. Thus, although something like U_N^{**} may approximate a statement of law, it’s not maximally perspicuous. There’s no reason to expect that laws can be perspicuously expressed in non-canonical terms, and in particular, in self-referential terms. Worries about circularity and/or redundancy would therefore seem ill-founded.

Finally, this primitivist stance suggests an optional but tempting conjecture regarding the relationship between statements of particular and general explanatory connections. According to this conjecture, for $[Fa]$ to be the ground of $[Ga]$ is just for $[Fa]$ to obtain, $[Ga]$ to obtain, and for a general explanatory law connecting F-facts with G-facts—e.g. $Fx \ll_x Gx$ —to obtain. On this view, for Jill’s maximizing happiness to ‘ground’ her having done something morally right is just for her to have maximized happiness, her to have done something right, and for any act of maximizing happiness to be thereby right. Given this identification, no problematic redundancy or circularity involving laws seems to arise.

6 Moral platonism

Although the nomic view is attractive, it’s not the only option for someone looking to vindicate the idea that moral principles are explanatory in role. To illustrate, we’ll consider an alternative: *moral platonism*.³⁵

Whereas the nomic view centers on the notion of a moral law, the platonist view of moral explanations centers on the distinction between particular-applying and kind-applying moral properties. The former are properties of *particulars* (i.e. datable, non-repeatable things), while latter are properties of *kinds* or *types* of things (i.e. timeless things with multiple instances). Initial motivation for this distinction comes from language: We not only say things like ‘Matti is good’ but also things like ‘Pleasure is good.’ On

³⁵ The platonist view is most thoroughly defended by Skarsaune (2015), who also discusses the linguistic motivation for it in more detail. For an earlier proposal along these lines, see Forrest (1986).

the face of it, whereas the predicate ‘is good’ applies to particular (Matti) in the former, it applies to a kind (pleasure) in the latter. The same is true of other moral terms. The view that moral predicates have genuinely kind-applying senses is further supported by facts about co-predication. More precisely, sentences that conjoin moral predicates with predicates that apply primarily, if not exclusively, to kinds (e.g. ‘Lying is wrong yet widespread’) are acceptable (Skarsaune 2015, 255).

Given that moral predicates like ‘good’ apply both to kinds and particulars in this way, we think the most plausible view is that this is a case of polysemy, as the predicate applies related but numerically distinct properties. This is supported by the zeugmatic nature of sentences that apply a single predicate to both particular-applying and kind-applying noun phrases—e.g. ‘MLK’s assassination was wrong, and so is lying’ feels a bit like a pun (cf. Skarsaune 2015, 258). Similarly, consider the following lists of things that are bad:

KINDS:	pain, suffering, unrequited love
PARTICULARS:	the pain you felt yesterday at 7pm, the war in Syria, Bill’s cheating on Hillary
MIXED:	pain, the pain you felt yesterday at 7pm, suffering, the war in Syria, unrequited love, Bill’s cheating on Hillary

We find KINDS and PARTICULARS considerably more natural than MIXED, despite the diversity of subjects within each. Indeed, the polysemy of moral predicates wouldn’t be surprising, as the distinction between particular- and kind-applying senses arises with many non-normative predicates, with certain sentences being ambiguous between the different readings. For example, ‘That’s a famous pen’ might be used to claim that a *particular* pen is famous (e.g. Einstein’s Waterman pen) or instead an instance of a famous *kind* of pen (e.g. a Fischer Space Pen).

What exactly is the relationship between kind- and particular-applying senses of predicates? In the non-normative case, kind-applying predicates are usually understood in terms of the particular-applying ones. Consider, for instance, the sentence ‘There was a dinosaur species that was 50 feet tall’ (Skarsaune 2015, 257). Intuitively, to say that a species was fifty feet tall is just to say that its adult members were (normally, approximately) 50 feet tall. But Skarsaune (2015, sect. 4) argues that the opposite relationship holds in the normative case: kind-applying predicates (and corresponding concepts) are basic and particular-applying ones derivative. This broadly linguistic thesis might then be taken as a guide to metaphysics, with particular-applying moral properties being analyzed in terms of kind-applying ones. This leads us to moral platonism.

According to platonism, moral terms like ‘good’, ‘wrong’, or ‘obligatory’ stand for two distinct but intimately connected properties. One property—call it ‘goodness_{kind}’—applies to kinds or types of things.³⁶ The other—call it ‘goodness_{part}’—applies to particulars. Particular-applying moral properties are then metaphysically analyzed (in a sense to be explained below) in terms of the more fundamental kind-applying ones: for a particular thing to be good_{part} is just for it to be a token of a good_{kind} kind; for a particular action to be wrong_{part} is just for it to be a token of a wrong_{kind} action-type; and so on. The basic moral properties are in this sense ‘second-order’—their relationship to their bearers is *instantiation*, in contrast to, say, necessitation (a la Armstrong 1983; cf. Murphy 2011).³⁷

Platonism straightforwardly accommodates data point (ii) and (iii)—concerning the role of descriptive facts and moral principles, respectively, in moral explanations. The Rossian principle of Fidelity, for example, can be formulated as follows (‘P’ for ‘platonism’):

PF_p: Lying is (pro tanto) wrong_{kind}

Assuming that to be a wrong_{part} action is just to be an action of a wrong_{kind} kind, PF_p guarantees that every instance of lying is wrong_{part}. For given PF_p, every particular lie will be an instance of a wrong_{kind} kind. This in turn will be explained by (a) the action’s being an instance of lying (a particular natural fact) and (b) lying being wrong_{kind} (a general moral fact). This fits the basic tripartite structure of explanation outlined in section 3. The platonist—like the nomicist—can thus adopt a version of the D&C strategy to make sense of data point (i), concerning Strong Supervenience.³⁸

³⁶ As Skarsaune (2015, 268) notes, platonism can be formulated in terms of properties (rather than kinds) having higher-order moral properties, with (e.g.) the property of being a lie being wrong-making. We’ll be treating claims about kinds/types and properties interchangeably; potential differences will be ignored.

³⁷ Rosen (2017: sect. 6.12) briefly considers an Armstrong-like view on which natural properties like *being a lie* stand in a ‘normative grounding’-relation (‘M’), rather than instantiation, to moral properties like *being wrong*. The fundamental moral facts have the following form: M(*being a lie*, *being wrong*). Rosen goes on to suggest that the M-relation is essentially such that if such a fact obtains and particular action is a lie, then its being a lie together with lying being M-related to wrongness metaphysically grounds the action’s being wrong (2017: 156). Thus, unlike platonism, Rosen’s proposal focuses exclusively on particular-applying moral properties and so fails to explain the application of moral terms to kinds. Worries also arise about the multiplication of grounding-relations, since the view posits two very different kinds of such relations, one of which (i.e. metaphysical) holds between facts while the other (i.e. normative) holds between properties.

³⁸ Platonism does not entail that every true sentence of the form ‘kind K is good’ states an explanatory principle. For instance, as Berker (p.c.) notes, if ‘Lying is pro tanto wrong’ is true, then ‘Lying-on-a-Tuesday is pro tanto wrong’ is plausibly also true. But the latter arguably doesn’t state a principle. While it’s hard to say exactly

6.1 Metaphysical analysis

It's worth clarifying the notion of metaphysical analysis being appealed to in describing how, on platonism, first- and second-order moral properties are related. It's sometimes thought, for example, that such analyses constitute a distinctive form of metaphysical explanation, stating what a fact or property 'consists in' (as opposed to, say, what they are 'grounded' in). Analyses are canonically expressed by sentences like 'To be F is to be G' and 'For something to be F is for it to be G'. For instance, one might accept Plato's proposal that to be virtuous is to have a well-ordered soul, or Aristotle's more recent proposal that to be human is to be a rational animal. More mundanely, one might think that to be a vixen is to be a female fox.

However, one complication is that claims of the form 'To be F is to be G' are capable of at least two readings. On what we'll call the *symmetric* reading, such claims express 'generalized identities' (cf. Linnebo 2014, Rayo 2014), or what Dorr (2016) calls 'identifications'. We can express generalized identities by introducing a 'no-difference' operator, \equiv , indexed by zero or more variables, which takes two open or closed sentences and yields another. As Corriea and Skiles (forthcoming) note, where p and q are open or closed sentences,

$$p \equiv q$$

can be read as 'For it to be the case that p (just) is for it to be the case that q '. Similarly, where F and G are monadic predicates,

$$Fx \equiv_x Gx$$

can be read as 'For a thing to be F (just) is for it to be G '. More generally, statements of the form

$$p \equiv_{x,y,\dots} q$$

can be read as 'For some things x,y,\dots to be such that p is for them to be such that q '. For example,

$$\text{Vixen}(x) \equiv_x \text{Female-fox}(x)$$

(For a thing to be a vixen is for it to be a female fox)

how the relevant distinction should be drawn, this is neither surprising nor particularly problematic—similar issues arise with respect to kind-predication and property inheritance more generally (cf. part-whole relations between particulars). See, e.g., Liebesman and Magidor (2017).

Like the familiar ‘objectual’ identity operator (e.g. Superman = Clark Kent), the generalized identity operator is reflexive, symmetric, and transitive. The above claim thus entails that:

$Female-fox(x) \equiv_x Vixen(x)$
(For a thing to be a female fox is for it to be a vixen)

In contrast, on the *asymmetric* reading of ‘To be F is to be G’, it expresses a corresponding general grounding claim, or ‘metaphysical law’ per above. So, for example:

$Female(x), Fox(x) \llcorner_x Vixen(x)$
(A thing’s being female together with its being a fox make it the case that it’s a vixen)

While the word ‘reduction’ has been used in (too) many different ways, symmetric readings of ‘To be F is to be G’ plausibly count as reductive insofar as they involve property identities (or something close). But they aren’t explanatory, given that metaphysical explanation is asymmetric. In contrast, while asymmetric readings are explanatory, they’re not straightforwardly reductive, since the facts related remain distinct. As the vixen example illustrates, there are often identificational truths in the vicinity of general grounding claims, and when there are the corresponding grounding claim is natural to accept. But the grounding claim neither is nor entails an identificational claim. Much the same can be said of sentences of the form ‘To be F is in part to be G’, such as ‘To be a vixen is in part to be a fox’. On one reading, they are general statements about what partially explains what. (Vixen-facts are partially explained by fox-facts.) On the other, they provide partial analyses. (For something to be a vixen is in part for it to be a fox.)³⁹

Note that statement of generalized identities, like statements of metaphysical laws, exhibit a kind of generality that entails but isn’t identical to universal generalization, standardly understood. Whether they are true or false, for instance, doesn’t depend on whether the left- or right-hand sides are satisfied in the actual world: even if there are no vixens, it’s still non-vacuously true that to be a vixen is to be a female fox.

³⁹ Thus, contra Schroeder (2005, 2007) and others, we don’t think metaphysical analyses constitute a distinctive, reductive form of metaphysical explanation. On one reading, the relevant claims are reductive but not explanatory; on the other, they are explanatory but not reductive (at least not on their own). We leave open the possibility of defining broader, non-identificational notions of reduction that supplement the grounding claim with, say, essentialist claims. See, e.g., Rydén’s (2018, ch. 3) development of Schroeder’s (2005, 2007) proposal.

Distinguishing metaphysical grounding from metaphysical analysis leaves room for what Rydén (2018, ch. 3) calls “metaphysically opaque grounding”. Simplifying slightly, this is the relation that holds when [P] is the ground of [Q] without being involved in the metaphysical analysis or essence of [Q]. This possibility conflicts with the commonly accepted idea that grounding guarantees a high degree of metaphysical intimacy. It has been claimed, for example, that if [P] is the ground of [Q], then [Q] is ‘nothing over and above’ [P]. This purported platitude is far from obvious, however, and arguably rests on a conflation of explanatory and identificational notions. (Compare: very few people think effects are ‘nothing over and above’ their causes.) Although one might think that general grounding claims, or metaphysical laws, are always accompanied by identificational facts, this is a substantive claim that requires argument. After all, it would seem to rule out various non-reductive metaphysical theses, including versions of metaethical non-naturalism and dualism about the mind. On the nomic view, for example, the non-naturalist will view natural facts as being “opaque” grounds of particular moral facts, explaining *why* but not *what it is* for particular actions to be right or wrong, good or bad, and so on.

Using the machinery just introduced, the platonist view of the relation between particular and general facts about (e.g.) wrongness can be formulated as follows:

x is wrong_{part} \equiv_x x is a token of a wrong_{kind} kind
 (For a particular thing to be wrong is for it to be a token of a kind that is wrong)

As noted above, generalized identity statements often guarantee certain general grounding claims. For example, given (PF_p), it’s plausible that:

Lying is wrong_{kind}, x is an instance of lying \ll_x x is wrong_{part}
 (Lying being wrong_{kind} and x being a lie together make it that x is wrong_{part})

This is a metaphysical law that, on the law-based view of grounding explanations, does explanatory work. Moral platonism is thus compatible with the law-based view of grounding explanation, though it’s a competitor to the nomic view of moral explanations. The difference between platonism and the nomic view concerns the relation between metaphysical laws and substantive moral principles. On platonism, the relevant moral principles are not metaphysical laws but facts involving kind-applying moral properties (e.g., the principle that lying is wrong takes the form [W(L)]). Metaphysical laws merely connect the moral principles plus the relevant natural facts with the explanandum in a transparent way. On the nomic view, by contrast, the

substantive moral principle is the law-like connection between natural and moral facts itself (e.g. $[x \text{ is lie} \llcorner_x x \text{ is wrong}]$).

7 Necessity

One question we've mostly ignored concerns the modal status of the fundamental moral principles. As noted above, the D&C strategy for explaining Strong Supervenience is only successful if the relevant principles obtain by metaphysical necessity. Non-naturalists will insist that those principles are not only morally fundamental but fundamental *simpliciter*. As a result, despite making sense of supervenience, they'll still be committed to there being *some* brute necessary 'connections' between the natural and non-natural.⁴⁰ This is hardly surprising—it's simply the bullet non-naturalists bite. Nonetheless, at least two more things can be said in closing.

To begin, although it's to be expected that non-naturalism will violate various 'Humean' metaphysical principles, the details matter. Take, for instance, the so-called 'Modest Humean' principle that "[c]ommitment to brute necessary connections between discontinuous properties counts significantly against a view" (McPherson 2012). Very different things might be meant by 'connections'—at a minimum, we should distinguish co-occurrence (iff), general 'making' (\llcorner_x) and instantiation (is). If the non-naturalist was forced to accept the necessary co-occurrence of particular moral and natural facts as brute, that would be quite bad. But they don't—as we've seen, they can make sense of such co-occurrence, and hence supervenience, by appeal to moral principles that are explanatory in role. Whatever worries remain won't be worries about *supervenience*—they'll be about something else.⁴¹

Another important issue concerns the sense in which the fundamental moral principles obtain 'of (metaphysical) necessity'. We think that, given certain assumptions about possible worlds, it's plausible for naturalists and non-naturalists alike to view the relevant principles as being *transcendental* in the sense of Fine (2005, ch. 9), rather than as (merely) necessary. The intuitive idea is that transcendental truths obtain *independently of*—rather than *in*—all worlds. Merely necessary truths, in contrast, depend on how things turn out in each world, coming out true every time. For example, seven's being prime is plausibly transcendental, while a truth like $P \vee \sim P$ is merely necessary—in some worlds it's true because P is while in others it's true because $\sim P$ is.

⁴⁰ We take Shafer-Landau's (2003) view that instances of moral properties are fully constituted by natural ones to be a form of naturalism, just as the view that instances of mental properties are fully constituted by physical properties is a form of physicalism (cf. Väyrynen 2017).

⁴¹ Needless to say, we're not suggesting that positing general explanatory principles is a metaphysical get-out-of-jail-free card. What's needed is independent support for positing them. And in the moral case, there is support for doing so.

The distinction between transcendental and merely necessary truths requires a more restrictive notion of necessity and truth-in-a-world than is standard. On Fine's alternative conception, a possible world is not to be viewed as a "totality of facts, or of how things might *be*, but [instead as] the totality of circumstances, or of how things might *turn out*" (2005, 325). While the totality of circumstances (roughly, substances and their properties and relations) varies between worlds, the transcendental facts form "the invariable framework within which the variation takes place" (325-6). Put in terms of logical form, while transcendental truths entail necessitated (or 'box'-ed) ones, they aren't themselves necessitated.

Like Fine, we doubt there's a need to choose between the more restrictive modal concepts and the more expansive ones. Different modal notions may be useful for different purposes. However, when it comes to understanding the modal status of many philosophical theses—including, but not limited to, fundamental moral principles—we think the more fine-grained concepts are most suitable. The analogies between morality and mathematics help illustrate this point. According to Ross (1930), for example, "[t]he moral order... is just as much part of the fundamental nature of the universe (and, we may add, of any possible universe in which there were moral agents at all) as is the spatial or numerical structure expressed in the axioms of geometry or arithmetic" (29-30). We doubt Ross is best interpreted as claiming that the relevant moral and mathematical facts are merely necessary—i.e. dependent on the circumstances in each world and coming out true every time. Rather, they are transcendental, obtaining independently of such circumstances and forming part of the framework in which the worlds play out.

Of course, whatever worries one might have about principles that are necessary in the broader, undifferentiated sense will likely carry over to principles viewed as transcendental. The appeal to transcendence is not intended to assuage such worries. The point is merely to shed additional light on how the view that such principles are necessary is best understood.

Bibliography

- Armstrong, D. M. 1983. *What is a Law of Nature?* New York: Cambridge University Press.
- Atiq, E. Manuscript. 'Supervenience, Repeatability, and Expressivism'.
- Audi, P. 2012. 'Grounding: Toward a Theory of the In-Virtue-Of Relation', *Journal of Philosophy* 109: 685-711.
- Berker, S. 2018a. 'The Unity of Grounding', *Mind* 127: 729-77.
- Berker, S. 2018b. 'The Explanatory Ambitions of Moral Principles', *Noûs*. doi:10.1111/nous.12246
- Berker, S. Forthcoming. 'Quasi-Dependence', in R. Shafer-Landau (ed.), *Oxford Studies in Metaethics*.
- Blackburn, S. 1984. *Spreading the Word*. Oxford: Clarendon Press.
- Bliss, R., and Trogdon, K. 'Metaphysical Grounding', in E.N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy*.
- Breckenridge, W., and Magidor, O. 2012. 'Arbitrary Reference', *Philosophical Studies* 158: 377-400.
- Carnap, R. 1963. 'Replies and systematic expositions', in P.A. Schilpp (ed.) *The Philosophy of Rudolf Carnap*. Open Court.
- Correria, F., and Schneider, B. 2012. 'Grounding: An opinionated introduction', in F. Correia and B. Schneider (eds.), *Metaphysical Grounding*. Cambridge: Cambridge University Press
- Dasgupta, S. 2014. 'On the Plurality of Grounds', *Philosophers' Imprint* 14: 1-28.
- DePaul, M. R. 1987. 'Supervenience and Moral Dependence', *Philosophical Studies* 51: 425-39.
- Dorr, C. 2016. 'To Be F Is To Be G', *Philosophical Perspectives* 30: 39-134.
- Dreier, J. 1992. 'The Supervenience Argument Against Moral Realism', *Southern Journal of Philosophy* 30: 13-38.
- Enoch, D. 2011. *Taking Morality Seriously*. Oxford: Oxford University Press.
- Enoch, D. Forthcoming. 'How Principles Ground', in R. Shafer-Landau (ed.), *Oxford Studies in Metaethics*.
- Faraci, D. 2017. 'Hybrid Non-Naturalism Does Not Meet the Supervenience Challenge', *Journal of Ethics and Social Philosophy* 12: 312-318.
- Fine, K. 2002. 'Varieties of Necessity', in T. S. Gendler and J. Hawthorne (eds.), *Conceivability and Possibility*. Oxford: Clarendon Press.
- Fine, K. 2005. *Modality and Tense*. Oxford: Oxford University Press.
- Fine, K. 2012. 'Guide to Ground', in F. Correia and B. Schneider (eds.), *Metaphysical Grounding*. Cambridge: Cambridge University Press.
- Fogal, D. 2016. 'Reasons, Reason, and Context' in E. Lord and B. Maguire (eds.), *Weighing Reasons*. Oxford: Oxford University Press.
- Forrest, P. 1986. 'Supervenience: The Grand-Property Hypothesis', *Australasian Journal of Philosophy* 66: 1-12.

- Glazier, M. 2016. 'Laws and the Completeness of the Fundamental', in M. Jago (ed.), *Reality Making*. Oxford: Oxford University Press.
- Hare, R. M. 1952. *The Language of Morals*. Oxford: Clarendon Press.
- Hattiangadi, A. 2018. 'Moral Supervenience', *Canadian Journal of Philosophy* 48: 592–615.
- Heathwood, C. 2012. 'Could Morality Have a Source?', *Journal of Ethics and Social Philosophy* 6: 1–19.
- Jenkins, C. S. 2008. 'Romeo, René, and the Reasons Why: What Explanation Is', *Proceedings of the Aristotelian Society* 108: 61–84.
- Kim, J. 1993. *Supervenience and Mind*. New York: Cambridge University Press.
- Kment, 2014. *Modality and Explanatory Reasoning*. Oxford: Oxford University Press.
- Kovacs, D. 2018. 'The Deflationary Theory of Ontological Dependence', *Philosophical Quarterly* 68: 481–502.
- Kovacs, D. Forthcoming. 'The Myth of the Myth of Supervenience', *Philosophical Studies*.
- Lange, M. 2018. 'What Would Normative Necessity Be?', *Journal of Philosophy* 115: 169–186.
- Leary, S. 2017. 'Non-naturalism and Normative Necessities', in R. Shafer-Landau (ed.), *Oxford Studies in Metaethics* 12. Oxford: Oxford University Press.
- Leary, S. ms. 'What is Moorean Non-naturalism?'
- Leuenberger, S. 2014. 'Grounding and necessity', *Inquiry*: 57: 151–174.
- Lewis, D. 1986. 'Causal Explanation', in D. Lewis (ed.), *Philosophical Papers*, vol. II. New York: Oxford University Press.
- Liebesman, D., and Magidor, O. 2017. 'Copredication and Property Inheritance', in *Philosophical Issues* 27: 131–66.
- Linnebo, Ø. 2014. "'Just is"-Statements as Generalized Identities', *Inquiry* 57: 466–82.
- Litland, J. 2013. 'On Some Counterexamples to the Transitivity of Grounding', *Essays in Philosophy* 14: 19–32.
- McNaughton, D., and Rawling, P. 2003. 'Descriptivism, Normativity and the Metaphysics of Reasons', *Aristotelian Society Supplementary Volume* 77: 23–45.
- McPherson, T. 2012. 'Ethical non-naturalism and the metaphysics of supervenience', *Oxford Studies in Metaethics* 7: 205–234.
- McPherson, T. 2015. 'Supervenience in Ethics', in E. N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy*.
- Miller, K. and J. Norton. 2017. 'Grounding: it's (probably) all in the head', *Philosophical Studies* 174(12): 3059–3081.
- Mitchell, C. 2017. 'Mixed up about mixed worlds? Understanding Blackburn's supervenience argument', *Philosophical Studies* 174: 2903–2925.

- Moore, G. E. 1942. 'A Reply to My Critics', in P. Schilpp (ed.), *The Philosophy of G.E. Moore*. Chicago: Opencourt.
- Murphy, M. C. 2011. *God and Moral Law*. New York: Oxford University Press.
- Rayo, A. 2013. *The Construction of Logical Space*. Oxford: Oxford University Press.
- Roberts, D. 2018. 'Why Believe in Normative Supervenience?', in R. Shafer-Landau (ed.), *Oxford Studies in Metaethics* 13. Oxford: Oxford University Press.
- Rosen, G. 2010. 'Metaphysical Dependence: Grounding and Reduction', in: *Modality: Metaphysics, Logic, and Epistemology*. Oxford University Press, pp. 109–36.
- Rosen, G., 2017. 'What is a Moral Law?', in R. Shafer-Landau (ed.), *Oxford Studies in Metaethics* 12. Oxford: Oxford University Press.
- Rosen, G. Forthcoming. 'Normative Necessity', in M. Dumitru (ed.), *Metaphysics, Meaning and Modality: Themes from Kit Fine*. Oxford University Press.
- Ross, W. D. 1930. *The Right and the Good*. Reprinted 2007. Oxford: Oxford University Press.
- Rydéhn, H. 2018. Ph.D. Dissertation, Uppsala University.
- Schroeder, M. 2005. 'Cudworth and Normative Explanations', *Journal of Ethics and Social Philosophy* 1(3): 1-27.
- Schroeder, M. 2007. *Slaves of the Passions*. Oxford: Oxford University Press.
- Schroeder, M. 2014. 'Price of Supervenience', in *Explaining the Reasons We Share: Explanation and Expression in Ethics, Volume 1*. Oxford: Oxford University Press.
- Schroeder, M. 2015. 'Being Realistic About Reasons, by T.M. Scanlon', *Australasian Journal of Philosophy* 93: 195-98.
- Schaffer, J. 2017a. 'The Ground Between the Gaps', *Philosophers' Imprint* 17(11): 1-26.
- Schaffer, J. 2017b. 'Laws for Metaphysical Explanation', *Philosophical Issues* 27: 302-321.
- Shafer-Landau, R. 2003. *Moral Realism*. Oxford: Oxford University Press.
- Shoemaker, S. 1987. 'Review of Blackburn's *Spreading the Word*', *Noûs* 21: 438-442.
- Skarsaune, K. O. 2015. 'How To Be a Moral Platonist', in R. Shafer-Landau (ed.), *Oxford Studies in Metaethics* 12. Oxford: Oxford University Press.
- Skiles, A. 2015. 'Against Grounding Necessitarianism'. *Erkenntnis* 80 (4):717-751
- Sturgeon, N. 2009. 'Doubts About the Supervenience Of The Ethical', in R. Shafer-Landau (ed.), *Oxford Studies in Metaethics* 4. Oxford: Oxford University Press.

- Toppinen, T. 2018. 'Non-Naturalism Gone Quasi: Explaining the Necessary Connections between the Normative and the Natural', in R. Shafer-Landau (ed.), *Oxford Studies in Metaethics* 13. Oxford: Oxford University Press.
- Väyrynen, P. 2013. 'Grounding and Normative Explanation', *Aristotelian Society Supplementary Volume* 87: 155–178.
- Väyrynen, P. 2017. 'The Supervenience Challenge to Non-Naturalism', in T. McPherson and D. Plunkett (eds.), *Routledge Handbook of Metaethics*: 170–84.
- 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.
- Woods, J. 2014. 'Logical Indefinites' *Logique et Analyse, New Essays on Logical Consequence* 57(227): 277–307.
- Woodward, J. 2014. 'Scientific Explanation', in E. Zalta (ed.), *The Stanford Encyclopedia of Philosophy*.