

promoting access to White Rose research papers



Universities of Leeds, Sheffield and York
<http://eprints.whiterose.ac.uk/>

This is an author produced version of a paper published in **Proceedings of the Aristotelian Society**.

White Rose Research Online URL for this paper:

<http://eprints.whiterose.ac.uk/76770/>

Paper:

Vayrynen, PP (2013) *Grounding and Normative Explanation*. Proceedings of the Aristotelian Society, 87 (1). 155 – 178.

<http://dx.doi.org/10.1111/j.1467-8349.2013.00224.x>

Grounding and Normative Explanation*

Pekka Väyrynen

University of Leeds

Forthcoming in *Proceedings of the Aristotelian Society*,
suppl. vol. (2013)

1. Introduction

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.¹ It is similarly very common for those who 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”. If we take this intuition seriously, then any brute and unexplainable normative facts there might be had better be quite special. We find this intuition not only in our thinking about how the

*Thanks to audiences at University of Leeds, University of York and the Ethics and Explanation conference at University of Nottingham for useful feedback on earlier versions of some of the material in this paper. Thanks to Chris Heathwood, David Plunkett and Debbie Roberts for comments that led to many revisions. And thanks to Christian Coons for planting the seed by forcefully pressing questions about normative explanation in conversations many years back.

¹By ‘normative’ I’ll henceforth mean ‘normative or evaluative’. In his contribution, Simon Kirchin raises concerns about this assimilation (Kirchin 2013). It doesn’t help that I won’t try to address the nature of normativity. I’ll work with the intuitive gloss that normativity involves standards of correctness which it is possible to violate and whose violation warrants some kind of criticism. I’ll also bracket such issues as whether the normativity of all that is normative can be explained by reference to reasons or any other single fundamental unit.

normative in general relates to the non-normative, but also in the background of many contexts of substantive normative inquiry, such as when normative ethicists seek to identify the right- and wrong-making features of actions to explain why certain actions are right and others are wrong.

This paper focuses on the “metanormative” issue of how these sorts of “normative explanations” work. Insofar as normative facts hold in virtue of non-normative facts, what in general is involved in this relationship? (A structural inquiry of this kind may not do more than provide background for answering questions about *which* non-normative factors ground *which* normative facts, and why they do so. Answering them will in addition require substantive normative assumptions.) I’ll focus in particular on whether recent work on “grounding” in metaphysics helps to shed light on how normative explanations work. A familiar line of thought about normative explanation – what I’ll call “the normative relevance argument” – implies that thinking of the explanatory relation at work in normative explanations as a grounding relation is of limited value in understanding how normative explanations work unless we take on controversial metanormative assumptions. The normative relevance argument can also be used to articulate some ways in which normative explanations might be distinctive from other domains where grounding has been appealed to. Much of this discussion will be fairly programmatic. But I hope the paper to kindle greater interest in how normative explanations work.

2. The Normative Relevance Argument

My interest in this paper is primarily in general claims like (1) and, derivatively, in such particular substantive instances as (2) and (3):

- (1) Actions have their normative properties in virtue of their non-normative properties.
- (2) Acts that maximize happiness are right in virtue of that fact.
- (3) This act is wrong in virtue of the fact that it was done with the sole intention of causing harm.

It is easy to generate variants of (1)-(3) by replacing ‘in virtue of’ with such related idioms as ‘because’, ‘depends on’, ‘is determined by’ and ‘makes’.² These claims are non-causal. For instance, (2) says not that maximizing happiness causes an action to be right, but rather that it makes an action right.

Claims like (1)-(3) are naturally interpreted as explanatory claims. Whatever else may go into the claim that maximizing happiness makes actions right, it implies that the fact that an act maximizes happiness plays some role in explaining why the act is right. Normative explanations like these are found all over the place, across various metanormative party lines.³ I’ll run my discussion largely in terms of cognitivist accounts of normative judgment, however. In this paper I am interested in what implications the structure of normative explanation has for the metaphysical structure of the normative. While I suspect that expressivist accounts of normative judgment (at least in their quasi-realist form) have resources to capture a lot of how normative explanations work, this might have little by way of metaphysical implications.⁴ It would be no surprise if the implications of how normative explanations work regarding the metaphysics of normativity depend on our account of normative thought and talk.

I’ll take explanation to be primarily a relation between facts.⁵ So what gets explained in normative explanation is a normative fact. I’ll assume that for a fact to be normative is for it to have a normative mode of presentation, and I want to leave open the possibility that one and the same fact could have both a

²I discuss these idioms as they appear in claims concerning the dependence of the normative on the natural in Väyrynen (2009a). This paper’s focus on the non-normative is different; for all I say here, the normative may be part of the natural.

³Here is just a small selection: Blackburn (1985, 37; 1988, 367-8), Dancy (1993, 79), Kim (1993, 225), Little (2000, 280), Smith (2000, 229) and Shafer-Landau (2003, 75). Some of these authors speak of “natural” rather than “non-normative” properties, but in most cases context makes clear that the author is using the two labels more or less interchangeably.

⁴Thanks to David Plunkett for discussion here. Also see Elstein (in progress).

⁵The relevant notion of explanation needn’t be thought of as one that is sensitive to various pragmatic considerations (such as background knowledge) or one under which some facts or propositions must have certain epistemic properties (such as being illuminating to beings in some cognitive predicament) to count as an explanation.

normative and a non-normative mode of presentation. (Perhaps to be right just is to maximize happiness.) This would be to think that facts are worldly items, individuated by their worldly constituents (such as objects and properties) and their manner of combination, independently of linguistic or conceptual guise. On this view, the fact that Sam is a bachelor and the fact that Sam is an unmarried male who is eligible to marry are the same fact. This is a plausible view about facts, but it is controversial all the same.⁶

Normative explanations like (1)-(3) can be read as *partial* or *full* explanations. But reading them as partial explanations wouldn't capture the dependence intuition. If the fact that an act maximizes happiness only partly explains why the act is right, completing the explanation might require some normative fact. If that normative fact can in turn be only partly explained in non-normative terms, its explanation may require appeal to some normative fact. So reading claims like (1)-(3) as partial explanations wouldn't allow us to capture the intuition that normative facts aren't in general brute and unexplainable. A robust dependence intuition seems to require reading at least some claims like (1)-(3) as claims about the full explanation of normative facts.

There is a familiar worry about the possibility of explaining *any* normative fact in purely non-normative terms. If I judge some action to be wrong and you ask me why it is wrong, I might say that it is wrong because it involves stealing, or because it is a case of promise breaking, or many other such things, depending on the act in question. No doubt this can be a sufficient explanation in the epistemic sense that someone who antecedently believes that stealing is wrong will conclude that the action is wrong when told that it involves stealing. But the following "normative relevance argument" seems to imply that there won't be full non-normative explanations of normative facts:

⁶In debates about grounding in metaphysics discussed below, this view of facts is rejected by Rosen (2010). The view raises special issues about how normative explanations are supposed to work under identity theories on which any fact that has a normative mode of presentation also has a non-normative mode of presentation (see e.g. McNaughton and Rawling 2003). I hope to discuss these issues elsewhere.

- (NR1) Some normative fact N is explained by a set of non-normative facts F.⁷ (Supposition)
- (NR2) The explanation of N by F presupposes that F has normative relevance: it is only if and because F is normatively relevant (in the right way) that it can explain N.⁸
- (NR3) The fact that F is normatively relevant (in the given way) is itself a normative fact.
- (NR4) This normative fact (that F is normatively relevant) cannot itself be explained by F.
- (NR5) So the explanation of N by F is incomplete; and to complete the explanation we need some further normative fact.
- (NR6) Since F and N were just schematic non-normative and normative facts, there can be no complete explanation of any normative fact that doesn't involve a further normative fact.⁹

To illustrate, suppose we offer the non-normative fact that x involves stealing as an explanation of the normative fact that x is wrong (NR1). The explanation that x is wrong because it involves stealing presupposes that the fact that x involves stealing is relevant to the wrongness of x (NR2). This presupposition would be satisfied if, for instance, stealing were a wrong-making feature of actions that have it; this is one way for stealing to be normatively relevant to the

⁷Throughout the paper, cases where I talk about F as a single fact will be the special case of a set with just one member.

⁸Here “in the right way” hides two qualifications. First, it is meant to require that F makes a difference specifically to that normative status which is instantiated in N (e.g. rightness or wrongness). Second, it is meant to allow for different modes of normative relevance. For instance, we might want to distinguish facts that are right- or wrong-making from facts that are normatively relevant by providing various sorts of background conditions (cf. Dancy 2004; Väyrynen 2006). Then the role of F in explaining N would depend on how it is normatively relevant.

⁹I take it that this argument sets out a familiar line of thought. I owe my particular formulation largely to Elstein (in progress). For related discussions, see e.g. Korsgaard (1996) and Schroeder (2005a).

wrongness of x . But the fact that stealing is a wrong-making feature is itself a normative fact, and the same goes for any other mode of normative relevance besides being wrong making (NR3). The normative relevance of the fact that x involves stealing cannot itself be explained just by the fact that x involves stealing (NR4). So the explanation that x is wrong because x involves stealing is incomplete. Completing the explanation requires some further normative fact, such as perhaps that stealing is a wrong-making feature (NR5).

Exactly the same reasoning seems to apply to any other pair of a normative fact and a set of non-normative facts, irrespective of whether the normative fact in question concerns reasons for action or value, rightness or wrongness *pro tanto* or overall and so on. So there can be no full explanation of any normative fact that doesn't involve a further normative fact.

One might read the normative relevance argument as a kind of regress argument. N isn't fully explained by F because F can explain N only if it is normatively relevant in the right way. What explains *that* normative fact? By the same argument, no set of facts expressed in purely non-normative terms does so; some further normative fact is needed. But what in turn explains *that* normative fact? And so on.

If the normative relevance argument is a regress argument, it has an important gap. It is one thing to say that F explains N only under a certain condition (here, that F is normatively relevant in the right way), another thing to say that this further condition is part of the explanation of N. So there is room to claim that what explains N is, simply, F, even if F can do that job only if it is normatively relevant in the right way.¹⁰ If the fact that F is normatively relevant needn't figure in the explanation of N by F, then no regress gets going. A further explanatory question might arise, of course: why *is* F normatively relevant? But closing one explanatory question and opening a further explanatory question isn't the same as regressing one and the same explanatory question. This might be thought to undermine NR2.

¹⁰Compare the discussion of the explanation of necessity in Hale (2002). Thanks to Ross Cameron for referring me to Hale's discussion.

But the normative relevance argument needn't be read as a regress argument. The feeling that F isn't by itself a sufficient explanation of N remains even if a regress of normative explanation is explicitly ruled out. Suppose you give me F, a reasonably rich set of facts expressed in purely non-normative terms, and tell me that because F is the case, I ought to ϕ . I could perfectly coherently reply "How so I should ϕ ?" (Perhaps I hold different normative views than you, or perhaps I am a normative nihilist who denies that there are any normative facts.) Simply repeating "because F" wouldn't be responsive to my query. Something more needs to be said to explain why I ought to ϕ . Adding further non-normative facts might well not help to remove the sense that citing purely non-normative facts as an explanation of a normative fact leaves the explanation hanging in the air. But saying something to the effect that F is a reason to ϕ or makes ϕ -ing valuable in some way, or that F would make it wrong not to ϕ , would be responsive to the query. So, in short, adding that F is normatively relevant in the appropriate way would be at least a first step towards securing the explanation, whatever form such an explanation in the end takes. I'll therefore suppose that NR2 is plausible at least on a reading that doesn't invite a regress of normative explanation.

So the normative relevance argument is at least initially plausible. But it seems to be in tension with the dependence intuition. So what now? Here is what we should ask: What is the explanatory relation at work here? Plugging different explanatory relations into the normative relevance argument might quite possibly lead to different assessments of its soundness or ramifications.

Below I'll focus on one particular proposal: the explanatory relation at work in the normative relevance argument is *grounding*. Here "grounding" is the standard term for the sort of non-causal dependence relation which is the subject of extensive recent discussion in metaphysics. I'll explain why appeal to grounding might promise to advance our understanding of how normative explanations work, but then argue that this promise rests on (ahem) shaky grounds.

3. Grounding and Normative Relevance

One major trend in recent metaphysics is that the grounding relation(s) expressible by saying that one thing holds *in virtue of*, *depends on* another or *is made the case by another* are to be understood as expressing robustly metaphysical rather than merely linguistic or semantic relations.¹¹ Grounding is typically introduced with examples such as these:

- (4) The proposition *Snow is white* is made true by the fact that snow is white. The truth-value of a proposition is determined by how the world is.
- (5) The singleton {Socrates} is grounded in Socrates. Non-empty sets depend for their existence in their members.
- (6) A glass is fragile in virtue of the arrangement of its constituent molecules (perhaps together with the laws of chemistry and physics). Dispositional features are grounded in categorical features.
- (7) The fact that Mary sees red obtains in virtue of some array of neurophysiological facts, perhaps together with certain laws. A phenomenal state is grounded in the neurophysiological state that realizes it.
- (8) If Jones means *addition* by '+', this semantic fact is grounded in some array of non-semantic facts.¹²

A standard point about examples like (4)-(8) is that the connection they feature doesn't seem analyzable in purely modal terms. Consider (5). Socrates and his singleton necessarily accompany one another, but Socrates isn't grounded in his

¹¹See e.g. Fine (2001; 2012), Schaffer (2009) and Rosen (2010). Interest in grounding often reflects a philosophical concern with how the reality is structured and which entities among those that exist are fundamental; see e.g. Fine (2001) and Schaffer (2009). For an alternative approach to such concerns using the notion of *metaphysical structure* instead, see Sider (2011).

¹²These sorts of examples can be found e.g. in Fine (1994; 2001; 2012), Schaffer (2009) and Rosen (2010), among many others.

singleton (Fine 1994). Or consider supervenience, the relation of necessary covariance. While the physicalist and the psychophysical parallelist in philosophy of mind will accept all the same claims regarding the supervenience of the mental on the physical, the parallelist won't accept the grounding claim in (7).¹³

Among typical illustrations of grounding are normative examples very much like (1)-(3). Here is Gideon Rosen, for example:

If an act is wrong, there must be some feature of the act that *makes* it wrong. Any given act may be wrong for several reasons, and some of these reasons may be more fundamental than others. A breach of promise may be wrong because it is a breach of trust, and a breach of trust may be wrong because it is prohibited by principles for social cooperation that no one could reasonably reject. (Rosen 2010, 110.)¹⁴

This explanatory idiom of “right- and wrong-making features” is widespread in contemporary moral theory but rarely analyzed.¹⁵ Appeal to grounding might be thought to help: claims like (1)-(3) might be thought to be neatly captured by saying that normative facts are grounded in non-normative facts.

A lot of work on grounding is devoted to specifying the properties of more robust relations of metaphysical determination or dependence.¹⁶ There is disagreement over these properties, and examples like (4)-(8) are sufficiently diverse to raise the question whether they exemplify any uniform formal structure.¹⁷ But across these disagreements many writers agree that the sorts of ‘because’ and ‘in virtue of’ claims they take to express grounding are a medium of *explanation*.¹⁸ For instance, (6) is naturally understood as conveying that the dispositional

¹³See e.g. Kim (1993, 167) and McLaughlin and Bennett (2005, §3.5).

¹⁴See also Schaffer (2009, 375), Audi (2012, 106-7) and Fine (2012, 37).

¹⁵Some exceptions include Dancy (2004), Strandberg (2008) and Väyrynen (2009a).

¹⁶See e.g. Schaffer (2009), Rosen (2010) and Fine (2012).

¹⁷For discussion, see e.g. Bennett (2011a) and Wilson (ms).

¹⁸See especially Fine (2001; 2012) and Audi (2012). See also Ruben (1990, ch. 7) for an early statement of a robustly metaphysical notion of non-causal explanation. Work on grounding should seek to clarify the sort(s) of non-causal explanations that ‘in virtue of’ and ‘because’ claims can be used to state.

features of things are, in some important sense, explained by their categorical features, and the same goes for truths and their truth-makers in the case of (4).

I'll largely bracket these details because my discussion of how grounding bears on normative explanation requires just that grounding is an explanatory notion, plus a couple of widely recognized formal properties of standard examples of grounding. These properties are that grounding is *asymmetric* (at least in its explanatory dimension) and that if P is grounded in some set of facts Q, then Q (or the conjunction of the propositions that correspond to the facts in Q) *entails* or *necessitates* (the proposition expressing) P.¹⁹ To say that grounding is asymmetric is to say that if Q grounds P, then it isn't the case that P grounds Q. While non-asymmetric relations such as supervenience run in both directions in cases like (5), plausibly the fact that {Socrates} exists is grounded in the fact that Socrates exists, but not *vice versa*. The dependence intuition says precisely that normative facts hold in virtue of non-normative facts, but not *vice versa*. The assumption that a grounded fact should be necessitated by its ground is more provisional. It is open to challenge from those who think that the relevant explanations needn't be that tight.²⁰

Saying that grounding involves asymmetric necessitation isn't sufficient to secure the kind of explanatory relation that grounding is supposed to be. Some purely modal connections involve these features as well. So for my purposes the specifically explanatory dimension of grounding remains a promissory note to be cashed in by specifying some further structural properties of grounding. There is disagreement about these properties.²¹ But I'll assume for argument's sake that there is some way to explicate what further properties are required to get

¹⁹Again cases where Q is a single fact will be the special case of a set with just one member.

²⁰Entailment/necessitation is controversial under certain interpretations of normative 'in virtue of' and 'making' claims; see e.g. Dancy (2004) and Zangwill (2008). In §4 I'll explain why giving up on entailment wouldn't help to deal with the dependence intuition.

²¹Grounding is plausibly both *non-monotonic* (see e.g. Rosen 2010, 116) and *hyperintensional* (see e.g. Schaffer 2009, 364). Both of these are also features of many explanatory relations. Properties that are more controversial in the case of grounding, but tangential to my purposes, are *irreflexivity* (see e.g. Jenkins 2010), *transitivity* (see e.g. Schaffer 2012) and *well-foundedness* (see e.g. Rosen 2010).

relations that are explanatory in the way grounding is meant to be. Skeptics about grounding will reject my project anyway.²²

Unfortunately already this very thin characterization of grounding leads to a problem for using grounding to illuminate how normative explanations work. Suppose that Q grounds P only if (the set of propositions expressing) Q entail (the proposition expressing) P.²³ Add to this the idea that grounding is an explanatory relation, and we get something like the following:

[A]ll that is properly implied by the statement of (metaphysical) ground itself is that there is no stricter or fuller account of that in virtue of which the explanandum holds. If there is a gap between the grounds and what is grounded, then it is not an explanatory gap. (Fine 2012, 39.)²⁴

If normative facts were grounded in non-normative facts in this way, it would underwrite structural grounding claims like that in (1). And although this wouldn't by itself suffice to identify any particular substantive normative facts of the form (2)-(3), it would secure the existence of some or other such normative facts.

The problem for using grounding to illuminate how normative explanations work is that the normative relevance argument directly entails that normative facts aren't *fully* grounded in non-normative facts in the kind of "strict" sense that Fine describes. The upshot of that argument is that any explanation of normative facts that cites only non-normative facts will be importantly incomplete. This implies that the strictest ground of any normative fact N will have to include some normative facts. Those who think that normative facts hold in virtue of non-normative facts cannot therefore mean by this that normative facts are fully grounded in non-normative facts in Fine's strict sense. The normative relevance argument allows them still to think that normative facts have a *partial*

²²For different types of skepticism about grounding, see e.g. Sider (2011), Daly (2012) and Wilson (ms).

²³This entailment is most naturally read as holding by metaphysical necessity.

²⁴It might be an epistemic gap instead; cf. note 5 above.

ground in non-normative facts in this strict sense.²⁵ But earlier we saw that even if claims like (1)-(3) are plausible as claims of partial ground or explanation, this won't help to capture any robust dependence intuition.

4. Three Reactions

I'll now describe three general reactions to the tension between the normative relevance argument and using grounding to illuminate how normative explanations work. These reactions will be available also to some candidates other than grounding for the explanatory relation in the normative relevance argument, including any relation that involves asymmetric necessitation. The lesson will be that we can understand the notion of explanation at work in the normative relevance argument in terms of grounding only if we take on certain controversial assumptions about the metaphysics of normative facts and properties.

One reaction is to try to weaken the grounding relation in some way to resolve the tension between the normative relevance argument and the claim that normative facts are fully grounded in non-normative facts. Getting rid of asymmetry seems inadvisable, so the move would have to be to revise the entailment/necessitation component of grounding. Caution is due here: separating grounding from any sort of entailment/necessitation might depart too much from other paradigm examples of grounding and make the notion of grounding too unconstrained to be useful. But suppose one says that the appropriate notion of grounding in the normative domain is a necessary connection of some type weaker than metaphysical necessity. The point I want to make is that it is by no means clear how appeal to such a relation of *normative grounding* would help to capture the dependence intuition.

Here is the worry. If P may be normatively grounded in Q without P being the case whenever Q is the case, what explains the difference between cases where {Q, P} and cases where {Q, ¬P}? The difference is explained either by some normative facts or by some non-normative facts. If the difference is

²⁵If P is grounded in several facts Q_1, \dots, Q_n taken collectively, then each Q_i is a *partial* ground of P and the Q s taken together are the *full* or *total* ground of P.

supposed to be explained by some non-normative facts, it seems that they can do so only if they are normatively relevant in the right way, and so the normative relevance argument gets going again. But if the difference is explained by some normative facts, then any set of non-normative facts would normatively ground some normative fact only given some further normative fact. This is what we get whether we think that normative grounding is relative to a background of normative laws or principles or that the fact that a set of non-normative facts F normatively grounds N is itself a normative fact. In either case we may again ask what grounds the normative facts in question.²⁶

It seems then that any set of non-normative facts would normatively ground some normative fact only given some further normative fact.²⁷ In general, if we say that a normative fact N is grounded and explained by a non-normative fact F but only given some suitable total background of further non-normative facts, this does nothing to touch the claims that (i) the explanation of N by F presupposes that F is normatively relevant in the right way, (ii) the fact that F is normatively relevant is itself a normative fact and (iii) this normative fact about F cannot itself be explained simply by citing F. So it is unclear how this reaction could hope to undermine the normative relevance argument.

Another reaction is to stand firm with grounding. One could maintain that the gap exposed by the normative relevance argument isn't a gap in the ground-

²⁶It is unclear in general whether the fact that Q grounds P itself requires a ground, and what that could be. For discussion, see e.g. Rosen (2010), Bennett (2011b) and Fine (2012).

²⁷Kit Fine proposes that there is a distinct and irreducible variety of *normative necessity* and a corresponding relation of normative grounding that licenses such claims as that the fact that a given act was right is normatively explained by the (non-normative) fact that it maximizes happiness (Fine 2012, 38-40; cf. Fine 2002). So Fine agrees that normative facts aren't fully grounded in non-normative facts as a matter of metaphysical necessity. I mention this here because Fine himself rejects the view that normative necessity is a matter of what is necessitated by the total state of the universe together with some relevant set of normative laws or principles (Fine 2002, 278). I don't have the space to discuss Fine's positive view here, but I suspect that the worry raised in the text can be raised also with respect to normative grounding as he understand it. The same applies to the more radically weaker accounts of normative explanation in Dancy (2004) and Zangwill (2008).

ing of normative facts by non-normative facts, but only a gap in our knowledge which can be eliminated by a fuller account of the grounds of normative facts. This would be to say that there is some set of facts F describable in purely non-normative terms and some fact N described in normative terms such that, if only we knew more about F, we would be able to see that it fully explains N.

The point can be made more precise in various ways depending on what exactly we say about grounding. Suppose, for instance, that we agree with Karen Bennett that grounding is a *superinternal* relation, “one such that the intrinsic nature of only *one* of the relata – or, better, one side of the relation – guarantees not only that the relation holds, but also that the other relatum(a) exists and has the intrinsic nature it does” (Bennett 2011b, 32).²⁸ This line of response to the normative relevance argument would then imply that the intrinsic nature of F is sufficient to guarantee that N holds, although we might not be able to see this if we don’t fully know the intrinsic nature of F. Different ways of explicating grounding generate variations on this reply. The common thread would be that there are some facts expressible in purely non-normative terms such that, given their existence, nothing else has to obtain for some normative fact to exist.

The broader metanormative commitments of this second reaction are an interesting issue. I am inclined to think that it isn’t open to *non-reductionism* about normative facts and properties. If N is a further fact over and above F, and in no way reducible to F, then not even full knowledge of the intrinsic nature of F would seem by itself to rule out the possibility that $\neg N$. For example, non-reductionists typically accept the supervenience of the normative on the non-normative, but the conjunction of F with the claim that there can be no N-difference without an F-difference doesn’t rule out the possibility that $\neg N$. Normative supervenience is compatible with normative nihilism, since if there are no Ns, it follows trivially that there can be no N-difference without an F-difference. Thus non-reductionists can get N out of F plus normative supervenience only given the further normative fact that some F somewhere is N.

²⁸This is in effect a restriction on Armstrong’s and Lewis’ notion of an internal relation as one whose holding supervenes on the intrinsic nature of the relata.

I suspect that the argument generalizes from supervenience to other necessary connections between distinct existences.

This tempts me to think that this second reaction to the normative relevance argument fits best with *reductionism* about the normative: normative facts are fully grounded in non-normative facts because facts expressible in normative terms are reducible to facts expressible in non-normative terms. Whether or not there is any systematic and general connection between grounding and reduction, various familiar reductionist idioms all suggest this sort of reaction.²⁹ Examples include such claims as that normative facts are “nothing over and above” non-normative facts or “consist in” in them, and that what it is for a certain normative fact to hold “just is” for certain non-normative facts to hold. This reaction fits with ideas according to which reduction can be thought of as a kind of *metaphysical analysis*.³⁰ Appeal to reduction would also fit with taking grounding as the explanatory relation as work in the normative relevance argument. It is usually thought that if A reduces to B, then B provides a particularly robust sort of explanation of A.

There would be much more to say about normative reduction than I have space for here, and than is usually said in metanormative theory.³¹ I’ll simply note that if understanding the notion of explanation at work in the normative relevance argument by appeal to grounding rests on reductionism, then it requires controversial assumptions about the metaphysics of normative facts and

²⁹The relation between grounding and reduction is discussed in Rosen (2010, 124-5) and Audi (2012, 110-1). Attribution of reductionism to this second reaction might not be apt if we thought that if P reduces to Q, then $P = Q$. There seems to be no reason to suppose that there should be any deep and systematic connection between grounding and identity. If reduction implied identity, this might also be thought to compromise the explanatory power of reduction (see e.g. McNaughton and Rawling 2003).

³⁰See e.g. Fine (2001) and Schroeder (2005b). Cf. King (1998) on the idea that *philosophical analysis* is a kind of property analysis. We needn’t think that reduction, understood as metaphysical analysis, must also yield a meaning analysis or semantic bridge principles, be knowable *a priori* or imply the eliminability of the reduced facts or properties.

³¹Notable exceptions include Railton (1986) and Schroeder (2005b; 2007, ch. 4). Oddie (2005, ch. 6) is a rigorous discussion of reduction by a non-reductionist.

properties. Whether normative reductionism is nonetheless the position to take in the end isn't something I'll try to settle here. But note that if the dependence intuition is best captured under normative reductionism, then a great many philosophers who are sympathetic to the intuition will have a hard time capturing it without giving up some other important commitment of theirs, such as their non-naturalism or constructivism about the normative. These philosophers would still owe us an account of what they mean when they say that normative facts hold in virtue of non-normative facts.

A third reaction is to say that at some point any normative explanation will hit the normative bedrock and can only appeal to some ungrounded normative fact that has no further explanation.³² This is to grant that the normative relevance argument is sound: no normative fact can be fully explained just by some non-normative facts. (This is to say that premise NR1 is plausible only as a claim about partial grounding or explanation, and false as a claim about total grounding or explanation.) Whether such "normative primitivism" is the position to take in the end isn't again something I'll try to settle here.

Clearly something more needs to be said, though. By this I don't mean that if F is bad, and this fact cannot be explained, then we should be able to explain why the badness of F cannot be explained.³³ But at least some "marks of the brute" seem reasonable to ask for. One such mark might be that if we are unable to give an explanation for a normative fact N after trying hard, then N is probably a brute and unexplainable normative fact. But it doesn't seem unreasonable to worry about taking our explanatory failure as grounds for thinking that N is so special as not to require or allow explanation. So what then would it take for a normative fact to count as brute?

Be that as it may, it is worth noting that although the primitivist reaction is most common among those who resists normative reduction, reductionism may be compatible with primitivism. For if the reductionist thesis itself (whether a

³²According to Russ Shafer-Landau, the moral realist must "say of the moral standards she favors that they just are correct – not in virtue of their being selected or created by anyone, but simply correct" (Shafer-Landau 2003, 46).

³³Thanks to Chris Heathwood for discussion here.

claim of property analysis, constitution or identity) turns out to have no further explanation, then even reductionists must accept ungrounded normative facts (Heathwood 2012).

I have described three reactions to the tension between the normative relevance argument and using grounding to illuminate how normative explanations work. I expressed skepticism about attempts to resolve the tension by weakening the grounding relation. But I didn't try to settle the plausibility of saying that normative facts are fully explained by, because reducible to, non-normative facts or conceding that some normative facts are ungrounded. Instead I'll propose that there is a distinction among normative explanations which constrains accounts of how normative explanations work.

5. Two Levels in Normative Explanation

Some kind of distinction among normative explanations seems afoot. It is one thing to say that some set of non-normative facts F is normatively relevant to N , another to say why F has its normative relevance. Saying that F is a reason for me to ϕ is one thing, saying in virtue of what F is a reason for me to ϕ is another. And saying that acts that involve stealing are wrong in virtue of that fact is one thing, saying why stealing is a wrong-making feature of acts is another. In the former cases we are after some (usually non-normative) feature that can be used to explain some normative feature, in the latter cases we are after that which bestows on the (non-normative) feature its normative relevance. If we take the dependence intuition seriously and so think that normative facts are at least typically not ungrounded, then it must be legitimate to ask what grounds the normative fact that F is normatively relevant to N , and this question isn't generally answered by simply citing F again.

This distinction can be captured by distinguishing two different explanatory questions. The first question is one that might be answered by claims such as *The fact that an experience would be painful is reason for you to avoid it*. This question concerns the *bearers* of (practical) normativity. (Under this heading we

can also ask which considerations are the *ultimate* bearers of normativity, and whether these are some sort of non-evaluative and non-normative facts, some sort of evaluative facts or some hybrid of the two.) So suppose that the painfulness of an experience is a reason for you to avoid it. A deeper explanatory question remains: In virtue of what does the fact that the experience is painful have the normativity of a reason? What is the *source* of the normativity that this consideration has? This source question is different from the bearers question. It asks in virtue of what the considerations that ultimately bear normativity – whichever they are, and indeed whatever normativity is – do so.³⁴

The fundamental point about this distinction isn't that only the source question is introduced in terms like 'in virtue of'. The bearer question can be addressed in such terms as well. We can say, for instance, that you would have a reason to avoid a certain experience in virtue of its being painful. Citing a bearer of normativity is at least normally relevant in explaining why something has a given normative property. But if the normative relevance argument is anything to go by, even any ultimate bearer of normativity can play a role in normative explanation only on the condition that it is normatively relevant in the requisite way. The source question asks in virtue of what a bearer meets that condition. So the fundamental point of the distinction is that normative explanations addressed to the bearer question concern a different normative role than those addressed to the source question.

How deep this distinction runs depends on whether to answer the source question is to identify some bearer of normativity. Some accounts of normativity will probably have this consequence, but this seems by no means obligatory. A Kantian might well say that the Categorical Imperative explains why certain considerations – whichever they are – are reasons to avoid certain actions but isn't itself a reason. (That might be "one thought too many", for instance.) A contractualist might well say that the fact that actions with certain features –

³⁴Different articulations of the source question can be found in Korsgaard (1996), Väyrynen (2006) and Chang (2009). The example of painfulness and the language of "bearers of normativity" is due to Chang (2009, 243).

whichever they are – are ruled out by any set of principles that no one could reasonably reject explains why those features make actions that have them wrong, but deny that such a fact about reasonable rejection is itself a wrong-maker.

So not all views imply that all explanations of why certain considerations – whichever they are – are bearers of normativity are themselves bearers of normativity by another name (Väyrynen 2009b, 101-2). Instead they can distinguish different levels or forms of normative relevance. One is to bear a certain kind of normativity (for instance, to have the normativity of a practical reason or to be a wrong-maker), the other is to provide a source of whatever normativity some consideration bears (for instance, to function as that in virtue of which some considerations are practical reasons or wrong-makers). Both are relevant to explaining why something has a given normative property, but in different ways. While “bearer explanations” may leave the source question unanswered, “source explanations” may get a grip in a particular context of normative explanation only once some relevant bearers of normativity have been identified.

In terms of this distinction, the upshot of the normative relevance argument is that any bearer explanation is at best a partial explanation of any normative fact if the bearer is itself non-normative. For instance, the claim that some set of non-normative facts F makes a particular action wrong hangs in the air as an explanation of why the action is wrong without some further account of where F gets the wrong-making force it bears. By a further iteration of the normative relevance argument, the normative fact that F bears wrong-making force cannot be explained in purely non-normative terms. The merit of the distinction is that it applies even to bearer explanations that cite evaluative considerations (such as that if something would be good for me, then in virtue of that fact I have a reason to go for it). Such bearer explanations answer the source question only if being evaluative in the way in question is something in virtue of which a consideration has the relevant kind of normativity. (In cases where that is so, the source question might be answerable by giving a bearer explanation. But that doesn’t mean that one sort of explanation generally collapses into the other.)

The distinction I have tried to articulate doesn’t require that we think of

normative explanations that address the bearer question and those that address the source question as featuring different kinds of explanation relation. The point is rather that complete explanations of why something has a given normative property must discharge two conceptually distinct explanatory burdens. This point remains even if on some substantive accounts of normative explanation these explanatory burdens aren't materially distinct because the source question can be discharged by citing some (ultimate) bearer of normativity.

More could be said about the bearer/source distinction than I have space for. It might, for instance, bear on how we understand theories in normative ethics. Many moral theories might be interpretable either as claims about what features of actions are (ultimate) bearers of rightness and wrongness or as claims about the source of the normativity that these features bear. In that case a proper understanding of a theory would require specifying which explanatory question it aims to answer. These issues must wait for another time. Instead I'll conclude by discussing how the distinction I propose bears on normative explanation.

6. Normative Explanations and Justification

Suppose accounts of how normative explanations work should accommodate the distinction between bearer and source explanations. Earlier I discussed the view that the explanatory relation in normative explanations like (1)-(3) is grounding. Looking at whether this proposal has resources to capture the bearer/source distinction helps us to see how it constrains accounts of normative explanation.

Return to the stock examples of grounding in (4)-(8). Suppose the singleton {Socrates} is grounded in Socrates. In virtue of what does Socrates ground {Socrates}? Or suppose that a glass is fragile in virtue of the arrangement of its constituent molecules. In virtue of what does that molecular structure ground the fact that the glass is fragile? These seem analogous to the source question in the normative domain. But if it is legitimate to ask in virtue of what Socrates has singleton-grounding force or a molecular structure has fragility-grounding force, the sensible answer is to say more about, respectively, Socrates and the

molecular structure. If there is an explanatory gap to close, that should close it.

The normative case is potentially different. One way to explain why is to say that normative explanations have a *justificatory* function. They aim to explain why things have properties such as rightness and wrongness. Most of us independently think that these properties somehow or other involve reasons – that is, considerations that *justify* actions. In that case we would expect normative explanations, too, to cite such considerations, and to be left hanging in the air otherwise.³⁵ Combining the bearer/source distinction with our earlier discussion of the normative relevance argument seems thus to have the upshot that saying more about (the intrinsic nature of) the bearer of normativity is likely to provide an adequate source explanation only if we take on controversial reductive assumptions about the metaphysics of normative facts and properties and only if citing these assumptions plays the appropriate sort of justificatory function. These conditions may be defensible, but they seem to make normative explanations disanalogous with grounding explanations in other domains where questions of normative justification don't arise.

This difference might be thought to reflect just the unsurprising fact that grounded facts have different features in different domains and this bears on what count as adequate grounding explanations in a given domain. But it may run deeper than that. The normative relevance argument implies that bearer explanations are typically plausible only as claims of partial ground or explanation.³⁶ It also implies that source explanations cannot be secured in purely non-normative terms without appeal to some normative facts. So the argument implies that normative explanations play their justificatory role only if they involve some further normative facts. This would sit ill with the idea that the explanatory relation in normative explanations is grounding. The characteristic ambition of appeals to grounding is to explain facts in one level by appeal to facts at some more fundamental level.

³⁵Thanks to conversations with Christian Coons and Daniel Elstein. Also see Elstein (in progress).

³⁶The exceptions would again be any cases where the source question can be answered by citing an ultimate bearer of normativity.

These remarks suggest constraints on reactions to the normative relevance argument. The reply that the normative relevance argument fails because normative facts can in fact be fully explained by citing the intrinsic natures of their non-normative grounds must tell us how citing those intrinsic natures is supposed to discharge the justificatory function of normative explanations. The version of this view according to which normative facts can be fully explained by non-normative facts thanks to a reduction must show that the proposed reduction captures the justificatory function of normative explanations. I am not saying that reductionists haven't recognized this constraint. Many have tried to show that their proposed reduction yields the right results regarding what is true about normative facts.³⁷ My suggestion is just that the constraint falls straight out of normative explanation. A fuller consideration of what goes into adequate normative explanations may also help sharpen other questions about normative reductions, including these: Is the reduction itself a normative fact capable of playing a role in normative justification? Is it a grounded or an ungrounded fact? What must be true of reduction for normative reductions to furnish normative explanations that fulfill also their justificatory function?

Now consider the view that some normative facts are ungrounded and have no further explanation. This view can seek partially to accommodate the dependence intuition. Suppose that the instantiation of normative properties requires the instantiation of non-normative properties as its basis. This need for a non-normative basis might not arise for ungrounded normative facts if these link the normative to the non-normative in some way that doesn't entail that any particular object instantiates any normative property.³⁸ Irrespective of whether these ungrounded normative link facts are thought to be reductive or non-reductive in character, this is an effective reaction to the normative relevance argument only insofar as grounded normative facts have their normativity in virtue of ungrounded ones in some way that fulfills the justificatory function of normative explanations. The primitivist should also tell us what it is about some normative

³⁷See e.g. Railton (1986) and Schroeder (2007).

³⁸See e.g. Shafer-Landau (2003) and Heathwood (2012).

facts that exempts them from explanation or justification. So again this function constrains a reaction to the normative relevance argument.

Much more remains to say about how different accounts of how normative explanations work interact with the justificatory function of these explanations, the pressure these accounts get from the normative relevance argument and so on. I must leave such questions for other occasions. But I hope that the issues I have been raising strike others, too, as interesting and worthwhile. They are fundamental in both metanormative theory regarding the nature and objects of normative thought and talk and substantive normative theory in ethics, aesthetics and elsewhere. How normative explanations work merits more attention.

References

- Audi, Paul (2012), “A Clarification and Defense of the Notion of Grounding”, in Correia and Schnieder, eds. (2012), 101-21.
- Bennett, Karen (2011a), “Construction Area (No Hard Hat Required)”, *Philosophical Studies* 154, 79-104.
- (2011b), “By Our Bootstraps”, *Philosophical Perspectives* 25, 27-41.
- Blackburn, Simon (1985), “Supervenience Revisited”, reprinted in Blackburn (1993), 130-48.
- (1988), “How to be an Ethical Antirealist”, reprinted in Blackburn (1993), 166-81.
- (1993), *Essays in Quasi-Realism* (Oxford: Oxford University Press).
- Chang, Ruth (2009), “Voluntarist Reasons and the Sources of Normativity”, in *Reasons for Action*, ed. D. Sobel and S. Wall (Cambridge: Cambridge University Press), 243-71.
- Correia, Fabrice and Schnieder, Benjamin, eds. (2012), *Metaphysical Grounding* (Cambridge: Cambridge University Press).
- Daly, Chris (2012), “Scepticism about Grounding”, in Correia and Schnieder, eds. (2012), 81-100.
- Dancy, Jonathan (1993), *Moral Reasons* (Oxford: Blackwell).
- (2004), *Ethics without Principles* (Oxford: Clarendon Press).
- Elstein, Daniel (in progress), “Property Identity and Normative Explanation”.
- Fine, Kit (1994), “Ontological Dependence”, *Proceedings of the Aristotelian Society* 95, 269-90.
- (2001), “The Question of Realism”, *Philosophers’ Imprint* 1.2. (hdl.handle.net/2027/spo.3521354.0001.002)
- (2002), “Varieties of Necessity”, in *Conceivability and Possibility*, ed. J. Hawthorne and T. S. Gendler (Oxford: Oxford University Press), 253-82.
- (2012), “Guide to Ground”, in Correia and Schnieder, eds. (2012), 37-80.
- Hale, Bob (2002), “The Source of Necessity”, *Philosophical Perspectives* 16, 299-319.

- Heathwood, Chris (2012), “Could Morality Have a Source?”, *Journal of Ethics & Social Philosophy* 6.2. www.jesp.org
- Jenkins, Carrie (2010), “Is Metaphysical Dependence Irreflexive?”, *The Monist* 94, 267-76.
- Kim, Jaegwon (1993), *Supervenience and the Mind* (Cambridge: Cambridge University Press).
- King, Jeffrey C. (1998), “What is a Philosophical Analysis?”, *Philosophical Studies* 90, 155-79.
- Kirchin, Simon (2013), “Evaluation, Normativity, and Grounding”, this volume.
- Korsgaard, Christine M. (1996), *The Sources of Normativity* (Cambridge: Cambridge University Press).
- Little, Margaret O. (2000), “Moral Generalities Revisited”, in *Moral Particularism*, ed. B. Hooker and M. Little (Oxford: Clarendon Press), 276-304.
- McLaughlin, Brian and Bennett, Karen (2005), “Supervenience”, *Stanford Encyclopedia of Philosophy (Fall 2005 Edition)*, ed. E. N. Zalta. plato.stanford.edu/archives/fall2005/entries/supervenience/
- McNaughton, David and Rawling, Piers (2003), “Descriptivism, Normativity and the Metaphysics of Reasons”, *Proceedings of the Aristotelian Society*, supp. vol. 77, 23-45.
- Oddie, Graham (2005), *Value, Reality, and Desire* (Oxford: Clarendon Press).
- Railton, Peter (1986), “Moral Realism”, *Philosophical Review* 95, 163-207.
- Rosen, Gideon (2010), “Metaphysical Dependence: Grounding and Reduction”, in *Modality: Metaphysics, Logic, and Epistemology*, ed. B. Hale and A. Hoffman (Oxford: Oxford University Press), 109-36.
- Ruben, David-Hillel (1990), *Explaining Explanation* (London: Routledge).
- Schaffer, Jonathan (2009), “On What Grounds What”, in *Metametaphysics: New Essays on the Foundations of Ontology*, ed. D. Chalmers, D. Manley and R. Wasserman (Oxford: Oxford University Press), 347-83.
- (2012), “Grounding, Transitivity, and Contrastivity”, in Correia and Schnieder, eds. (2012), 122-38.
- Schroeder, Mark (2005a), “Cudworth and Normative Explanations”, *Journal of*

- Ethics & Social Philosophy* 1.3. www.jesp.org
- (2005b), “Realism and Reduction: the Quest for Robustness”, *Philosophers’ Imprint* 5.1. hdl.handle.net/2027/spo.3521354.0005.001
- (2007), *Slaves of Passions* (Oxford: Clarendon Press.)
- Shafer-Landau, Russ (2003), *Moral Realism* (Oxford: Clarendon Press.)
- Sider, Theodore (2011), *Writing the Book of the World* (Oxford: Clarendon Press).
- Smith, Michael (2000), “Does the Evaluative Supervene on the Natural?” Reprinted in his *Ethics and the A Priori* (Cambridge: Cambridge University Press, 2004), 208-33.
- Strandberg, Caj (2008), “Particularism and Supervenience”, in *Oxford Studies in Metaethics*, vol. 3, ed. Russ Shafer-Landau (Oxford: Clarendon Press), 129-58.
- Väyrynen, Pekka (2006), “Moral Generalism: Enjoy in Moderation”, *Ethics* 116, 707-41.
- (2009a), “Normative Appeals to the Natural”, *Philosophy and Phenomenological Research* 79, 279-314.
- (2009b), “A Theory of Hedged Moral Principles”, *Oxford Studies in Metaethics*, vol. 4, ed. Russ Shafer-Landau (Oxford: Clarendon Press), 91-132.
- Wilson, Jessica (ms), “No Work for a Theory of Grounding” (unpublished).
- Zangwill, Nick (2008), “Moral Dependence”, in *Oxford Studies in Metaethics*, vol. 3, ed. Russ Shafer-Landau (Oxford: Clarendon Press), 109-27.