
White Rose Research Online

eprints@whiterose.ac.uk
https://eprints.whiterose.ac.uk
Universities of Leeds, Sheffield and York

Deposited via The University of Leeds.

White Rose Research Online URL for this paper:
https://eprints.whiterose.ac.uk/id/eprint/203750/

Version: Accepted Version

Book Section:
Väyrynen, P. (2024) Explanatory Pluralism in Normative Ethics. In: Timmons, M., (ed.) Oxford Studies in Normative Ethics Volume 14. Oxford University Press, pp. 138-161. ISBN: 9780198930785.

https://doi.org/10.1093/oso/9780198930785.003.0008

© Pekka Väyrynen 2024. This is an author produced version of a book chapter published in Oxford Studies in Normative Ethics Volume 14. Uploaded in accordance with the publisher's self-archiving policy.

Reuse
Items deposited in White Rose Research Online are protected by copyright, with all rights reserved unless indicated otherwise. They may be downloaded and/or printed for private study, or other acts as permitted by national copyright laws. The publisher or other rights holders may allow further reproduction and re-use of the full text version. This is indicated by the licence information on the White Rose Research Online record for the item.

Takedown
If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.





Explanatory Pluralism in Normative Ethics

Pekka Väyrynen
University of Leeds

1. Introduction

Theories in normative ethics don't aim merely to list which things are good, bad, right, and wrong, and to guide us accordingly. They aim also to say *why* the right actions are right and the wrong ones are wrong. Consequentialists and Kantians may agree that a given action is wrong, but disagree on whether it's wrong *because* it fails to maximize value or *because* it violates the dignity of persons. Explanation is important also in everyday ethical thought. Most people accept a general moral presumption against lying but think that lying isn't always wrong. Just when lying is wrong depends on why it's wrong. For lying isn't wrong when the *reasons why* it's wrong are outweighed or fail to hold altogether (Korsgaard 2007: 577). All this goes also for other areas of normative inquiry: we want to know what *makes* things good for us, beautiful, politically legitimate, and so on. Explanation is thus a core enterprise of first-order normative inquiry. Just witness titles like "Why Is Death Bad?", "What Makes Killing Wrong?" and "What Is Wrong with Lying?"[1]

Moral theorists have said rather less about what kind of explanations their answers to these questions are meant to state. This is probably justified. Scientists typically work with recognized explanatory practices of their fields and are probably justified in leaving it for philosophers of science to systematize and account for the nature of these explanations. But many accounts of scientific explanation have failed to account for some genuine explanations that practicing scientists give by privileging certain types of explanations.[2] I'll argue that existing theories of normative explanation run into the same problem. Most assume *normative explanatory monism*, the view that all explanations of particular normative facts which are appropriate in first-order normative inquiry work the same way. I'll instead defend *normative explanatory pluralism*, the view that a particular normative explanandum may have more than one type of correct complete explanation. I'll focus on normative ethics, but much of what I'll say will apply also to many other areas of normative inquiry.

This is a preprint of a paper forthcoming in *Oxford Studies in Normative Ethics* 14 (2024).

1 See, respectively: Brueckner and Fischer (1986) and Kamm (1993: Ch. 1); Sinnott-Armstrong and Miller (2013); and Faulkner (2007) and Korsgaard (2007).

2 A notorious example: simple Deductive-Nomological models of explanation (Hempel and Oppenheim 1948) rule out inductive and probabilistic explanations and explanations that require no reference to laws.

2. ‘Normative’, ‘Explanation’, and ‘Pluralism’

Explanation involves a thing being explained (the “explanandum”) and its explainer (the “explanans”) related in a certain kind of way (the “explanation relation”). The term ‘explanation’ may refer just to the explanans, as when we say that the explanation of Q is that R. Or it may refer to the whole relationship: when we ask whether an explanation is correct, we don’t merely ask whether the putative explanans is true. Normative explanations are some type of non-causal explanations of why something normative is the case. Both ‘normative’ and ‘explanation’ are said in bewilderingly many ways. Here I can only quickly fix some ideas that aim to reflect plausible assumptions about what normative ethicists are up to in their explanatory enterprise. This will help to identify the relevant form of pluralism.

As for ‘normative’, it’s enough to work with paradigmatic examples. Normative facts and claims include the fact or claim that an action is morally wrong, that it’s what we morally ought to do, that some outcome is morally good, that some institutional arrangement is unjust, and the like. If explanatory pluralism is plausible in the case of paradigmatic moral explananda, it won’t be less plausible for broader classes of normative fact.[3] Here I won’t fuss much about just how ‘moral’ and ‘normative’ relate. Statements of moral explanations, such as ‘Lying is wrong because it treats people as mere means’ are typically classified as a kind of normative claim. I won’t worry about which other complex statements with moral terms in them, such as ‘Lying is wrong or grass is green’, are normative claims.

Many moral facts and claims are particular, such as that it was wrong that I lied to you about where I was last night. Most particular moral facts don’t hold brutely. But it’s controversial whether all such facts have explanations. Perhaps “nothing explains why what it’s like to be in agony is a reason for me to avoid future agony” (Berker 2019: 931). Other moral facts are general in various ways. Most of us accept a general moral presumption against lying, Rossian pluralists think that lying is *pro tanto* wrong, and ‘Lying is wrong’ may also be used to make a generic moral claim akin to ‘Birds fly’. Some general moral facts have explanations: intermediate moral generalizations are explained by whatever they derive from. Like laws of nature, fundamental moral principles raise complications that I’ll bracket here. Whether and how they can be explained depends on their nature – a controversial matter.

As for ‘explanation’, I’ll make standard assumptions: explanation is asymmetric, irreflexive, factive, non-monotonic, and hyperintensional. A more substantive assumption that seems to be implicit in much of moral inquiry is that moral explanation is moderately

3 For more complications with defining ‘normative’ than you ever wanted to face, see Finlay (2019). The notion of a normative fact may here be understood in a metaphysically deflationary fashion.

objective. One aspect of this objectivity is that a moral theorist's proposed explanation of why (say) lying is wrong isn't fundamentally a speech act of explaining with respect to the question of why lying is wrong. The explanation itself isn't subject to general pragmatic influences on communicative acts even if the explanation can only be put across to others through such acts. We can say all this and think that explanations have the property of being the *content* of an answer to a why-question. There can be facts about what possible answers to a given question are correct even if we haven't given an answer, or even asked the question. Explanations needn't themselves *be* things out in the world to count as objective.

A more complex issue is whether the truth-conditions of moral explanations are worldly. Moral theorists don't seem to be making an epistemic or pragmatic claim when proposing that an act is wrong because it fails to maximize (actual or expected) value, or because it treats others as mere means, or the like. Suppose I broke a promise to you and this was wrong. Perhaps it was wrong because it failed to maximize utility, or for some other reason. Whatever explanation is correct, its truth doesn't depend on whether it increases its audience's understanding, answers to their interests, or the like. A true explanation becomes no less true if a given audience lacks the conceptual resources to understand it or it's old news to them (Lewis 1986: 226-28). Explanations in moral theory don't have this kind of deeply audience-relative truth-conditions. Weaker epistemic conditions may not compromise a significant degree of objectivity. For instance, perhaps genuine explanations "have the potential to increase the understanding of rational agents with proper training" (Kovacs 2020: 1672-73). I won't try to settle whether having such potential is a condition on the truth of a moral explanation, or an extra good.

This fix of ideas aims to reflect plausible assumptions about explanations that moral theorists typically are after. Here's a first pass on pluralism about that class of explanations:

> *Normative Explanatory Pluralism* (NEP): For some normative explananda N, 'N because R' states a correct complete explanation of type A and 'N because Q' states a correct complete explanation of type B (where N is a normative explanandum and A ≠ B).[4]

How many moral facts have more than one type of explanation is a substantive question (see §4). Normative explanatory monism is false if at least one does. An interesting kind of

4 On the moderately objective conception, these explanations may exist without having been given. I also won't require that R ≠ Q. Given certain further assumptions, R and Q may sometimes coincide; see §6.

pluralism requires some principled distinction among types of explanation; I'll offer one in §5. My aims require the relevant types to be explanations that are appropriate in first-order moral inquiry. NEP therefore says more than that non-causal explanations in ethics, mathematics, and science may lack a unified account, more than that moral facts may also have explanations that are of interest to metaphysics but not moral theory, and more than that they may have both moderately objective and deeply epistemic or pragmatic explanations.

NEP also makes a stronger claim than Kenneth Walden's explanatory pluralism (Walden 2016). Walden notes that one and the same moral fact may have more than one correct explanation when it conceals multiple explananda corresponding to different contrast spaces. My valet drops a tray. I thrash him with a blackjack. I did wrong. That it was wrong to *thrash* him (instead of some milder rebuke) is a different explanandum than that thrashing him was *wrong* (rather than permissible). 'Why was it wrong to *thrash* him?' and 'Why was thrashing him *wrong*?' are different questions, answered by different explanations (Walden 2016: 192). Since NEP ranges over normative explananda, it applies only once a normative explanandum has been individuated, contrastively or otherwise.

My defense of NEP is simple. First, I'll show that moral inquiry traffics in at least two types of explanation of particular moral facts. Contrary to monists who think that moral explanation takes only one or the other as its fundamental or standard form, each is perfectly appropriate. To say this, we must accept something along the lines of NEP. Second, I'll show that NEP can offer principled responses to some central objections.

3. Two Normative Explanatory Monisms

Talk of some facts holding in virtue of other facts, making something the case, and making other facts obtain is often naturally interpreted as *generative:* it concerns mechanisms that *give rise* to phenomena. Normative explanations of this type might be characterized like so:

> *Generation:* R explains a normative explanandum N insofar as (and because) R is a plurality of facts that bring N about.[5]

Explanation by *Generation* provides a fine-grained account of the specific factors that give rise to the explanandum, through some form of asymmetric metaphysical determination. The 'because' signals that this is what gives generative explanations their explanatory power. If

5 R may be a single fact in partial explanations and in limiting cases of complete explanation.

you accept both *Generation* as a model for normative explanation and normative explanatory monism, you think that *Generation* specifies the only genuine type of normative explanation.

Some moral explanations seem generative. Ordinary moral thought often focuses on particular instances of wrongful behavior and what it was that made them wrong. The wrongness of many acts is said to result directly from how they inflict physical suffering on someone or substantially curtail their autonomy. The wrong in a marital transgression is often explained by a detailed account of how the transgression happened. We say things like: what made it wrong for government officials to award a lucrative contract to Pestilence (a company run by a substantial donor to the party of government) without comparative tender is that this was corrupt. One intuitive model of how certain personal relationships make a moral difference is that certain duties are directly grounded in them: those relationships generate duties where none existed before (cf. Stroud 2024).

Generative explanation figures in more self-consciously theoretical projects as well. A Rossian pluralist explanation of why my failing to turn up when I had promised to do so was wrong overall might seek to specify not only that I thereby failed to fulfill the weighted balance of *prima facie* duties but also that my failing to turn up makes it the case that I failed in my *prima facie* duty of fidelity and didn't contribute to the fulfillment of any weightier *prima facie* duty. According to act-utilitarianism, rightness and wrongness always result from facts about the net balances of happiness and misery that each alternative action available to the agent would have brought about.

You think *Generation* has instances if you think there are *grounding explanations* of particular moral facts, where 'grounding' denotes a kind of asymmetric metaphysical determination. People who think this disagree over details. One dispute is whether a complete grounding explanation of why an act is wrong needs to cite only certain non-moral facts or also a moral principle. The former view might be supported by a grounding-explanatory analogue of bare causal explanations that require no reference to laws (cf. Skiles 2015: 742-44) or putative explanatory redundancy of moral principles (Berker 2019). On the latter (perhaps more common) view, my examples of generative moral explanation are elliptical and tacitly appeal to general moral principles (Rosen 2017; Fogal and Risberg 2020).[6] Another dispute is whether moral explanation involves the same form of grounding whereby a vase is fragile because of its microphysical structure, or a distinct type of normative

6 An alternative to the view that moral principles, together with non-moral facts, generate the moral fact that's being explained is that they govern the grounding hierarchy while sitting outside of it (Bader 2017).

grounding (Fine 2012; Berker 2018; Enoch 2019). But all these views are typically presented in a way that suggests that all explanations of particular moral facts work by *Generation*.

A clear instance of this form of normative explanatory monism is Selim Berker's claim that debates between consequentialists and their opponents cannot be made proper sense of if 'right-making characteristics' is understood as picking out conceptual priority, epistemic priority, supervenience, logical entailment, counterfactual dependence, or identity, but that they can be made sense of if 'making' is understood as picking out metaphysical grounding (Berker 2018: 738-43). Berker suggests the same is true of debates about the "reasons first" approach to normativity (Berker 2018: 743-46). Based on such examples, he draws a monistic inference: first-order investigations of normative notions have all along (self-consciously or not) been "fundamentally in the business of proposing (and assessing, and establishing) various grounding claims" (Berker 2018: 743).[7]

Does *Generation* identify the *only* genuine type of normative explanation? Some putative explanations draw connections that don't seem generative. One type looks like this:

> *Unification:* R explains a normative explanandum N insofar as (and because) R includes some entity E more general than N such that N and many other phenomena, seemingly disparate from N, can be derived from E.[8]

Unifying explanation is often said to work by locating the explanandum in a broader pattern. Think of this as a convenient shorthand. Not every locating in a broader pattern yields a true instance of *Unification*. One slightly more precise measure of this feature is the extent to which a theory can present a large number of phenomena (generality) as consequences of a small set of explanantia, from which they can be derived via relatively few and similar patterns of derivation (simplicity) that impose various kinds of restrictions on the instances of those patterns (cohesion). The 'because' signals that derivability of the explanandum from something more general by employing few widely applicable patterns of derivation is what gives unifying explanations their explanatory power. Many unificationists think that the relevant general entities are statements of laws or principles and the patterns of derivation are argument schemas (Kitcher 1989). But these items needn't be fundamentally linguistic or

7 Sachs (2018: 20) also endorses grounding explanation monism about explanations of particular moral facts.

8 See e.g. Kitcher (1989), Schurz (1999), and Strevens (2004) for unificationist views of scientific explanation. *Unification* roughly states a guiding idea shared by these philosophers; details vary a lot. Kovacs (2020) and Baron and Norton (2021) extend unificationism to metaphysical explanation. Such extension doesn't require thinking that grounding relations are obsolete (Stamatiadis-Bréhier forthcoming).

propositional. They might be more basic facts or general properties and correspondingly worldly connections among phenomena. Either way, unifying explanation as such requires no reference to bottom-up metaphysical determination. Unifying derivations are often treated as subsumptive, but there could be non-subsumptive patterns of derivation (cf. Cullity 2018: Ch. 4). If you accept both *Unification* as a model of normative explanation and explanatory monism, you think that *Unification* specifies the only genuine kind of normative explanation.

Unificationism is often defended by saying that noticing the connections and common patterns that tie phenomena together increases understanding (Kitcher 1989: 431-32). But there's no deep tension between *Unification* and my earlier suggestion that explanations of importance to moral theory may have no essential tie to understanding. First, again, potential to increase understanding might be an extra good, not a truth-condition on explanation. Second, the explanatory value of unification needn't depend on its potential to increase understanding. Patterns of derivation that contribute to unification might have explanatory power insofar as they track relations or patterns that obtain independently of our inferential practices.[9] W. D. Ross sought to explain facts about what we morally ought to do by identifying the smallest number of distinct *prima facie* duties, and argued that our duties of justice, beneficence, and self-improvement can be unified under a single *prima facie* duty to promote intrinsic value (Ross 1930: 26-27). Ross seems not to take the explanatory value of this result to consist primarily in its potential to increase understanding. The same goes for Garrett Cullity's view that the full normative and evaluative content of morality can be derived from three foundational norms governing concern for others' welfare, respect for their self-expression, and cooperation, using a limited set of patterns of derivation (Cullity 2018).[10] In general we think that specific rights or duties shouldn't be treated as *sui generis* unless we really have no other option. The more of them can be derived from more general rights or duties, the fewer normatively basic moral requirements we need to accept and justify as authoritative. This explanatory virtue of unification isn't primarily epistemic.

Literature in ethics is full of accounts that are plausibly construed as instances of *Unification*. It's standard to treat various specific duties and rights as derivative. For example, why is exploitation at least *pro tanto* wrong even when consensual and mutually beneficial? Its wrongness isn't fundamental. Perhaps an exploitative transaction is wrong when and

9 Kitcher's original unificationist framework treats causal relations as reflections of our inferential practices (Kitcher 1989: 436-37). The point in the text is that *Unification* doesn't require a corresponding form of antirealism about non-causal relations between phenomena. Cf. Stamatiadis-Bréhier (forthcoming: 13).

10 Cullity describes derivations of norms in non-epistemic terms (Cullity 2018: 13).

because the exploiter gains an unfairly greater share of the benefits of social cooperation than the exploited; its wrongness derives from unfairness or distributive injustice (Roemer 1996; Mayer 2007). Or perhaps exploitation is wrong when and because it violates our moral obligation not to extract excessive benefits from people who cannot, or cannot reasonably, refuse our offers (Valdman 2009). Or perhaps it's wrong when and because one party benefits from taking advantage of the other party's vulnerability in a way that involves violating proper respect for others (Wood 1995: 150-51) or dominating others via subordination (Vrousalis 2013). Each view seeks to explain the wrong of exploitation as an instance of some broader wrong, typically by subsuming it under more general principles or values.

Another example is the duty to provide treatment for kidney failure. Those who accept such a duty and the corresponding right to treatment typically explain them as an instance of a collective duty for people within a nation-state to provide health care to people in need and a corresponding right for people in need to receive the necessary help. The more specific duty qualifies as an instance of the more general duty because organ failure is a catastrophic treat to people's health and well-being. It might even be thought to derive from even more general rights, such as that people in need have a right to the material resources they require to lead decent lives. Certain body parts, including kidneys, might be thought sufficiently like material resources that the sick at least sometimes to have a right that the healthy give them some of the body parts they need to lead a decent life (Fabre 2003). Debates in bioethics tend to focus on how the more general duty is best fulfilled in the special case of kidney failure (Sterri 2021 is one helpful summary). But whichever form a duty to provide treatment for kidney failure takes, its standard explanation derives it from some more general duties.

Explanations of various injunctions of justice provide further examples. Suppose justice requires the eradication of poverty and that this in turn requires distributive policies or institutions that are, in Ingrid Robeyns's term, "limitarian": in the world as it is, no one should have more than a certain upper threshold of resources, such as income, wealth, or natural resources (Robeyns 2022). People disagree over the proper basis of the limitarian threshold, but in most proposals the basis is some more general value or principle. For instance, perhaps the reason why the poor should be assisted through redistribution of excess wealth is that the poor have a quality of life that's below a sufficientarian threshold, or that the needs of the poor have weighted priority, or that excess wealth contributes to undermining political equality (Robeyns 2022 surveys the debate). This proposal is also an example of another pattern: accounting for various more specific rights or duties by purporting to derive

them from core democratic values. Consider the right of those affected by a governmental decision to participate in making that decision. James Wilson's account of this participatory right derives it from the core democratic values of autonomy and equality. My autonomy will be objectionably disregarded if I'm subject to direction or implication of my will without authority over its terms. The moral equality of persons implies that in sufficiently dense social interactions, each participant has an equal entitlement that the decisions of other participants don't objectionably disregard her autonomy in this way. (Wilson 2022: 171-74.)

In many of these examples, the official explanandum isn't a particular moral fact. Robeyns is primarily interested in accounting for a limitarian principle, Wilson for the "all-affected principle" of democracy. But if these can be explained on the basis of some more general principle or value, then so can their particular instances. Witness ordinary moral thought. Why was it wrong for Nixon to lie about his administration's involvement in the break-ins to the DNC headquarters? We might say "because he lied". But most of us don't regard lying as brutely wrong. So we might reflect further, not just on this case but other cases as well. Through reflection, we might come to think that Nixon acted wrongly because of how such lies undermine civic trust. Even if we took civic trust as explanatory bedrock, we would have sought to explain the wrongness of Nixon's conduct by showing it to be an instance of a more general wrong. Still further reflection might lead us to conclude that undermining civic trust in turn is wrong because of how it involves lack of respect for others, or because civic trust is part of what democracy is.

These first-order moral investigations may not be couched in explicitly unificationist terms. For example, often we're told little about what general patterns of derivation are employed. But it's natural to reconstruct them as instances of *Unification*. *Unification* implies that at least strong forms of moral particularism are false. But so do these investigations. In each, some more specific moral phenomenon is said to be explained on the basis of some more general moral value, ideal, or principle. A further unifying feature is that explanation of particular moral facts and explanation of the mid-level moral generalities of which they're instances have the same basic form. The measures that contribute to unification (generality, simplicity, and cohesion) distinguish how unifying explanation works from how generative explanation works. The more a set of derivations can unify phenomena, the more we explain by less. Generative explanation as such has no essential tie to these measures.

Explicit appeals to unification have been rare in discussions of moral explanation. One exception is Derek Baker's account of moral explanation as a form of *unifying*

generalization (Baker 2021). Whether the explanandum is why a particular lie is wrong or the mid-level generality that lying is wrong, its wrongness can be explained by subsumption under some more general property it shares with some larger class of wrong acts. If this more general property is (say) treating others as mere means, it might unify a particular lie not just with other lies that are wrong but also with further, superficially different ways of treating people as mere means that are also wrong. Based in part on the kinds of examples I've highlighted, Baker draws a monistic inference: "[T]he standard form of a normative explanation is a unifying generalization, one which explains by presenting a more particular normative injunction as a special case of a more general injunction" (Baker 2021: 57).

4. Why Pluralism?

Generation and *Unification* set out conceptions of explanation that have various attractions. As philosophy goes, we can expect some people to insist that moral theorists' interest in the most fundamental right- and wrong-making characteristics of actions means that at bottom they should seek only generative explanations, and others to insist that moral theory most fundamentally traffics in unifying explanations of varying levels of generality. But the rest of us don't need to choose, at least if we bracket general concerns about each conception that aren't specific to explanation in ethics. I'll now defend pluralism over monism.

When I introduced NEP in §2, I said that normative explanatory monism is false if at least one normative explanandum has more than one type of explanation that's appropriate in moral inquiry. This is sufficient but not necessary. Monism is false already if some particular moral facts have generative explanations while others have unifying explanations. Above we saw that each type is routinely pursued in first-order moral inquiry.[11] It's easy to find more examples of each type. That should alleviate worries about moving from evidence that moral theorists *routinely* seek more than one type of explanations of particular moral facts to conclusions about what moral explanations are *genuine* or *appropriate*. If either generative or unifying moral explanations weren't genuine, the theorists who have sought them would have to have fundamentally misconceived their project. Such imperialism has rarely turned out well for theories of scientific explanation. Explanation in ethics looks no different.

This case against monism is stronger still if there are more than two types of moral explanation. Marc Lange argues that many mathematical explanations involve "explanation

11 It's not always clear whether a proposed explanation is best understood as working by *Generation*, *Unification*, or some third way. Often the issue is presentational. For example, sometimes a right is said to be "grounded" in certain interests, but is explained by derivation from those interests' moral importance.

by constraint", which explains a phenomenon by appeal to its absence being an impossibility of a robust kind (Lange 2016: 16). Now notice that if it's metaphysically impossible for any act whose maxim cannot consistently be willed as a universal law to be morally permissible, this impossibility constrains the moral status of all acts, regardless of their specific non-moral details or what more general kinds they fall under. A good question for future work is whether such explanations are appropriately characterized as explanations by constraint.

NEP thus makes a stronger claim than just the negation of monism. NEP says that at least in some cases, *one and the same* normative explanandum has more than one type of correct complete explanation. It may be that, as a matter of fact, some moral facts commonly receive one type of explanation while others commonly receive another type. But I see no principled basis for saying that *these* moral facts have only a generative explanation, *those* have only a unifying explanation, and the two classes are exhaustive and mutually exclusive. For example, the distinction couldn't go by way of the moral property in the explanandum. My examples in §3 include both generative explanations of why some act was wrong and unifying explanations of why some other act was wrong.

It's a substantive matter just which particular moral facts have more than one type of explanation. Some may have just one type of explanation, or even none. For example, not just any derivation of an explanandum, or its locating in a broader pattern, yields a true instance of *Unification*. (Different accounts of unification may set that threshold differently.) It's an open question whether all particular moral facts have derivations that reach the degree of generality, simplicity, and cohesion that a genuine unifying explanation must have.

I hope that a slightly stylized example will make it plausible that many normative explananda have more than one type of explanation, provided that some moral generalities required for unifying explanation obtain. Adapting an example from Selim Berker, let:

W = [She acted wrongly in telling Al]

S = [She could have done something else instead of telling Al that would have brought about more overall happiness]

L = [She could have lied instead of telling Al, her lying would have brought about 100 overall units of happiness, and her telling Al brought about 20 overall units of happiness]

Assuming that W obtains in virtue of S and S in virtue of L, “it is very natural to hold that W also obtains in virtue of L” (Berker 2018: 752). If L explains W, it does so by *Generation*. But now consider:

W′ = [She acted wrongly in letting Bea die]

D = [She could have done something else instead of letting Bea die that would have brought about more overall happiness]

A = [She could have prevented Bea’s death by making Colin lose a limb, her making Colin lose a limb would have brought about -40 overall units of happiness, and her letting Bea die brought about -100 overall units of happiness]

By parallel reasoning, if A explains W′, it does so by *Generation*. Each of W and W′ thus has at least two generative explanations: W by S and by L, W′ by D and A. L and A specify the mechanisms that give rise to W and W′ in a more fine-grained way than S and D. Nothing in the general structure of generative explanation seems to privilege one explanation over the other. But their finer-grained differences may be of moral interest. For example, in the former case even the wrong action is good while in the latter case even the right action is bad.

A unifying explanation of W and W′ abstracts away from the level of non-moral detail found in L and A. Including that level of detail would require more complexity in the patterns of derivation, and thus a lower degree of unification. A unifying explanation may also abstract away from some of the detail found in S and D. For instance, it doesn’t really matter in a unifying explanation that one act is telling Al and the other is letting Bea die. The most unifying patterns of derivation operate on the comparative structure: a different available option than what the agent chose would have brought about more overall happiness.

The lesson of the example isn’t restricted to act-utilitarianism. For example, Kantian unifying explanations would similarly abstract away from the details of the particular maxims whose universalizability is what helps to unify the morally permissible acts, while generative explanations may well cite various more specific features incorporated in the maxims of permissible actions. Richard Yetter Chappell argues more generally that what’s “criterial” for rightness needn’t be what “most fundamentally makes” an act right; the features that a theory posits as having “ground-level moral significance” can be more specific than its general criteria for right or wrong action (Chappell 2021: 427). Provided that criterial principles of right and wrong also are explanatory of particular moral facts, this view predicts that such

facts will often have both unifying and generative explanations. On this view, maximizing act-consequentialism doesn't imply that *failing to maximize value* is the fundamental wrong-making feature. If so, it won't figure in the fundamental generative explanations of why wrong acts are wrong; those will feature gratuitous harm, lying, inequality, and the like. But explanations that cite criterial principles are systematically going to instantiate more unifying patterns than explanations that cite specific right-making features. Indeed, omitting details of particular phenomena sometimes enables a greater degree of unification than would otherwise be possible (Strevens 2004: 156-57).

How might normative explanatory monists respond? Since adding and subtracting finer-grained detail impacts generative and unifying explanations differently, prospects seem dim for arguing that one sort of explanations somehow supplant or reduce to the other. The response I've encountered most often from moral theorists is that generative moral explanations that cite only non-moral features are elliptical and tacitly appeal to general moral principles (together, perhaps, with other assumptions that connect more specific non-moral features to those that figure in the principles), and this collapses them into unification. In §6 I'll argue that even if we grant the former point, the latter doesn't follow.

There are further reasons to accept NEP. One is that NEP does a good job accounting for phenomena associated with the property that explanations have of being contents of answers to why-questions. First, one and the same why-interrogative can often be used to ask different why-questions. It's familiar that why-interrogatives can be used to ask for causes, grounds, goals or purposes, and more. In some contexts 'Why was it wrong of Kim to lie?' may be used to ask what brings about this particular moral fact. Then only grounds that give rise to it qualify as possible answers. In other contexts, the same interrogative may be used to ask how this moral fact is like other wrongs. Then only unifying generalizations count as possible answers. Different possible answers to explanatory why-questions correspond to different explanations. But we needn't think that the explanandum embedded in these why-questions is different. Second, we can ask not only questions like 'Why is lying wrong?' but also such more specific questions as 'What makes it wrong to lie?' and 'Is the wrongness of lying like other wrongs, and how?' These aren't why-questions on their face, but an answer will typically be an answer also to the question of why lying is wrong. This also suggests that one and the same normative explanandum may have different types of explanation.[12]

12 Thanks to Daniel Fogal for helping me to get clearer on this point in another context.

NEP also avoids a common objection to monism about scientific explanation: any monism requires all scientific explanations to have a certain feature, but no feature seems to be a necessary feature of all scientific explanations.[13] If unification isn't necessary for explanation, this is a problem for monism premised on *Unification*. If not all explanations work by showcasing in detail what gave rise to the explanandum, this is a problem for monism premised on *Generation*. Neither is a problem for NEP. More tricky is sufficiency. One way to defend monism is to argue that all but one putative type of explanation are insufficient for explanation. Many worry that unification isn't enough for explanation. Perhaps unification doesn't guarantee asymmetry (Gijbers 2007). Or perhaps what matters to explanation is giving a more profound account of a phenomenon or stability under some relatively wide range of counterfactual interventions, whether or not either contributes to unification with other phenomena (Ruben 2012: 191-92; Woodward 2003). Elsewhere I raise similar worries about *Generation:* at least generic metaphysical grounding of a particular normative fact neither suffices for its normative explanation nor always provides a normatively deeper account of why it obtains (Väyrynen 2021b). Here NEP and monism face the same question: just what *is* sufficient for each type of explanation?

Finally, NEP is parallel to a common form of pluralism about scientific explanation. Scientific explanations of a phenomenon may vary even within a given discipline when in service of different research projects. For instance, a biological trait such as scarlet ibis feather color may have a physiological explanation (one that cites genetic or other mechanisms whereby the scarlet ibis produces feathers with carotenoid pigmentation), an explanation in terms of evolutionary game theory (one focused on the status of certain scarlet ibises in terms of their feather of this color), an explanation focused on the role of selection (such as an ecological model of the role of this coloration in mate attraction), and more (Potochnik 2020). These explanations focus on different sets of causal factors or patterns and isolate them using models that are idealized and simplified in correspondingly different ways (Potochnik 2015; Mantzavinos 2016). In recent work, Alex Stamatiadis-Bréhier extends this more general kind of pluralism to moral explanation. On his view, whether P is an "appropriate explanans" for Q is relative to the aims and goals of the research project in which the explanation "P explains Q" figures (Stamatiadis-Bréhier ms). Grounding explanations can serve the aim of showcasing in great detail the mechanisms that brought about a particular moral fact. Unification, by contrast, can serve the aims of prediction and

13 Rice and Rohwer (2020) provide a useful summary of various arguments to this effect.

understanding with respect to the same explanandum. An explanation can serve one of these aims without serving the other. I won't here compare Stamatiadis-Bréhier's view with mine; our main concerns are different.[14] Explanations that I take to be of chief importance to moral theorists may have no essential link to prediction or understanding.

This more specific form of explanatory pluralism brings up an important point: NEP is best understood as a schema that can be filled out in different ways depending on how we type normative explanations. If we interpret Stamatiadis-Bréhier as distinguishing types of explanation relative to the aims or goals of research projects, then his view is a special case of NEP. A central challenge to NEP in any form is whether the types of explanation it ranges over can be individuated in a suitably principled way. Different aspects of this challenge raise three objections to NEP. A key to my replies is that *Generation* and *Unification* pick out explanatory connections that are distinct in kind even when they're co-instantiated.

5. Too Many Types of Explanation?

One worry about NEP is *overgeneration* of correct explanations. One version of the worry is that NEP allows questionable projects of moral inquiry to carve out a type of normative explanation such that explanations of that type can be correct.[15] Suppose I think that justice requires prioritizing the interests of white people. I can pose myself an explanatory question: Why is it just to prioritize the interests of white people? My inquiry has a goal: answering explanatory why-questions about what justice requires. No doubt I could rustle up an answer to the question of why prioritizing the interests of white people is just. Why couldn't there be a type of explanation that can be a correct and complete answer to my explanatory question?

To this we should respond that explanation is factive with respect to the explanandum: 'R explains N' entails N. Since justice doesn't require prioritizing the interests of white people, the explanandum is merely putative. Other questionable projects are ruled out by factivity with respect to the explanans: 'R explains N' entails R. Some explanations fail not because their explananda don't obtain, but because their explanantia are or presuppose something false. This would be an atheist's diagnosis of why divine command explanations of moral facts fail, for example.

The overgeneration worry doesn't die this easily. Suppose "consequentializers" are right that any non-consequentialist moral theory has a deontically equivalent consequentialist

14 Stamatiadis-Bréhier uses pluralism to respond to Berker's (2019) explanatory redundancy argument: even if moral principles are redundant in grounding explanations, they're not redundant in unifying explanations.

15 Thanks to Sam Baron and Alex Stamatiadis-Bréhier for versions of this objection.

counterpart. Insofar as these theories offer different explanations of their shared deontic verdicts, they should still be distinct theories. One might then worry that there's one type of explanation that correctly explains all moral facts in terms of (say) the contravention of the Categorical Imperative and another that correctly explains the same moral facts in terms of the contravention of some consequentialist principle. This would weaken explanation's role in the debate over which moral theory is correct. That would be a problem for NEP.

This objection also dissolves on closer inspection. For there to be a worry, the Kantian and consequentialist explanations must differ in more than their explanantia. The objection only shows that 'Lying is wrong *because* [fill in a Kantian first-order story]' and 'Lying is wrong *because* [fill in a consequentialist story]' may each have a true explanans. But the truth of 'R explains N' requires more than that R and N both obtain. It might *be* that any wrong act is, by necessity, one that God commands us to refrain from, without this being a *reason why* the wrong acts are wrong. As the objection stands, nothing suggests that Kantian and consequentialist explanations are distinct types of explanation in the relevant sense.

Distinct types of explanation must pick out or highlight different types of relation or connection between the explanans and the explanandum, whatever these relata may be. Again different explanations of scarlet ibis feather color pick out different causal mechanisms and so aren't in competition. Conjoining them may yield what's in some sense a more comprehensive overall account of this biological trait. But that composite isn't a distinct type of explanation; no new explanatory connection is introduced. *Generation* and *Unification* similarly invoke different types of connection: more or less fine-grained metaphysical structures that give rise to phenomena in one case, fairly abstract structures of derivation that tie phenomena together in the other. Kantians and consequentialists may offer rival generative explanations of moral facts. They may also offer rival unifying explanations of the same facts. Since each type draws different explanatory connections, we don't compare a Kantian unifying explanation against a consequentialist generative explanation. Within each type, which explanation is true depends on substantive first-order matters. Nor do we expect that one and the same moral fact might have a correct Kantian unifying explanation but a correct consequentialist generative explanation, provided that their substantive commitments conflict. Kantian and consequentialist accounts thus don't constitute further distinct types of explanation *qua* Kantian and consequentialist explanations, simply in virtue of their different substantive contents. NEP thus survives these overgeneration worries.

6. One Type to Rule Them All?

I just offered a necessary condition for when types of explanation are distinct. My argument requires further that generative and unifying explanations, specifically, come out as distinct. Otherwise the distinction between them fails to motivate NEP. But there's a worry that these *collapse*.[16] Consider stylized examples of each type:

(GR) Alex's lying was wrong because it caused Hamid to suffer needlessly.

(UN) Alex's lying was wrong because it contravened a general moral principle not to cause needless suffering.

Even if (UN) is correct complete explanation of why Alex's lying was wrong, it's far from obvious that (GR) is. Many grounding theorists accept the following constraint:

Grounding Sufficiency: A plurality of facts Γ is a full ground for Q only if Γ is sufficient, on its own, for Q.

For instance, the parts ground the whole only if the parts, on their own, are enough for the whole. But many moral theorists think that the fact cited in (GR) isn't sufficient on its own for Alex's lying to be wrong. If there's no reason to prevent needless suffering in general, the fact that a lie would cause needless suffering wouldn't help ground a duty not to lie (or lying's being wrong) in any particular case. Instead, (GR) tacitly appeals to a general principle, such as that cited in (UN). On this way of thinking, a complete grounding explanation looks like this:

(GR-COMP) Alex's lying was wrong because it caused Sam to suffer needlessly and there's a general moral principle not to cause needless suffering.

But (GR-COMP) seems like a unifying explanation. If any particular moral fact will have a general moral fact among its full grounds, *Generation* and *Unification* won't be distinct.

The objection is twofold: (A) a complete grounding explanation of a particular mora fact cites a general principle as a partial ground and, therefore, (B) it will also be a unifying explanation. (A) is controversial. For instance, Selim Berker argues that moral principles are

16 I first got this worry from Sam Baron, and subsequently from various others.

redundant or superfluous in grounding explanations of particular moral facts (Berker 2019). But even if (A) is true, (B) doesn't follow. As we saw in §4, *Generation* and *Unification* pull in different directions. All else equal, abstracting away from finer-grained detail enables a greater degree of unification (a greater number of moral explananda can be derived via fewer patterns of derivation), but makes a generative explanation less informative. An explanation of a particular moral fact that instantiates a highly unifying pattern would often not do so if it were enriched with more specific detail concerning the factors that give rise to that moral fact. Unifying and generative explanations have different properties *qua* such.

What may be true of some cases of the form (GR-COMP) is that they co-instantiate generative and unifying explanatory relations. The point remains that even if a complete grounding explanation of an explanandum cites the same general principle as its unifying explanation, the principle's unifying work is *incidental* to its role in grounding explanations, and *vice versa*. Explanatory incidentality is an independently recognized phenomenon. Marc Lange writes: "That a distinctively mathematical explanation happens to cite facts about the explanandum's causes does not mean that it works by virtue of describing the explanandum's causes" (Lange 2016: 20). Similarly for the claim that all correct causal explanations meet a condition of counterfactual dependence: some explanations support counterfactuals that relate changes in their explanantia to changes in their explananda only incidentally (Khalifa *et al*. 2020: 1446). For example, some unifying explanations whose function is to highlight unity in apparent diversity would realize this function even if they didn't support the relevant kind of counterfactuals (Khalifa *et al*. 2020: 1451-57). My response to the collapse worry is parallel.

A unifying explanation works by providing a certain kind of derivation of the explanandum. Even if its explanans cites facts that also ground the explanandum-fact, that's incidental to its unifying-explanatory power. A grounding explanation works by describing the mechanisms that give rise to the explanandum. Even if its explanans also unifies, that's incidental to its grounding-explanatory power. Nothing about grounding explanation as such requires its explanans to have any significant degree of generality. A universal principle can be highly specific and still play the role it plays in (GR-COMP). But given how we've unpacked *Unification*, there's some level of specificity below which a principle won't contribute to unifying explanations. So a principle's power to unify is distinct from its power (if any) to bring about its instances. Since *Generation* and *Unification* pick out different explanatory relations, they have distinct truth-conditions. The collapse worry may have more sophisticated variants, of course. One is that we can get unification by abstracting from the

right-making characteristics cited in a set of generative moral explanations. Another is that generative explanations can be run within a unificationist framework so long as the former's explanans contains reference to something that a moral theory can abstract away from, revealing a pattern that can be included among the inferential resources that the theory uses to give moral explanations.[17] I suspect my response can be adapted to deal with these worries.[18]

7. Not a Genuine Pluralism?

A final worry is *deflation: Generation* and *Unification* mark distinct types of explanation, but these are different enough phenomena not to constitute an interesting form of pluralism. A common version of the worry says that *Generation* identifies a metaphysical phenomenon whereas *Unification* identifies an epistemological phenomenon, but it's no surprise whatsoever that these are distinct. For example, some think that unification is best understood as making sense of why the explanandum obtains, but it's plain enough that providing a fine-grained account of what makes moral facts obtain (which looks metaphysical) can come apart from making sense of why those facts obtain (which looks epistemological).

However, *Generation* and *Unification* might not be aptly contrasted as "metaphysical" versus "epistemological". If "metaphysical" explanation is something that proceeds by citing distinctions that are of importance to metaphysics, such explanations aren't automatically of importance to first-order normative inquiry. In any case, unificationism is one account of metaphysical explanation in this sense (Kovacs 2020; Baron and Norton 2021). Perhaps explanation is instead "metaphysical" (better: objective) when its truth doesn't depend on its potential to provide cognitive benefits such as increasing understanding. But as noted in §3, one interpretation of unification allows this kind of objectivity. NEP can thus claim to state a pluralism within a shared subject matter. The types of explanation that normative inquiry may seek have distinct truth-conditions. What holds them together is that they draw different types of explanatory connection between a normative explanandum and its explanans.

More still may hold *Generation* and *Unification* together. NEP allows that generative and unifying explanations of moral facts may share interesting necessary conditions.

17 For an analogue of the latter idea, see Strevens (2004; 2008) on how to integrate causal explanations into a unificationist framework. Thanks to David McElhoes for pressing me on this worry.

18 Jamie Dreier raised a different collapse worry. It seems we can chain these explanations. For example, if A grounding-explains B and B unifying-explains C, it seems that A explains C. Chainability would be surprising if the two were deeply distinct. But even if generative and unifying explanations are sometimes chainable, I doubt this holds in general. For example, it's unclear that chaining will systematically transmit the respective explanatory virtues of generative and unifying explanations. If it doesn't, in what sense does chaining yield an explanation? For related discussions concerning the chainability of different types of explanation, see e.g. Berker (2018: 751-55), Duguid (2021), Väyrynen (2021b), and Lange (2023).

Elsewhere I defend a necessary condition on normative explanation that would apply to both types (Väyrynen 2021a-b). That condition would also help NEP answer a general problem for explanatory pluralism.[19] As a mere necessary condition, it won't suffice to unify normative explanations into a single type. But it's only if they meet it that different types of connection between a normative fact and its explanans qualify as *normative explanations*.

I've defended normative explanatory pluralism: one and the same normative explanandum may have more than one type of correct complete explanation – even when we focus just on explanations of importance to moral theory. I've argued that explanatory pluralism in normative ethics is supported by what goes on in actual moral inquiry, parallels an independently plausible form of pluralism about scientific explanation, avoids common problems with explanatory monism, and can offer principled responses to central objections.[20]

19 My answer is implicitly structured after Chris Pincock's (2018: 42) dilemma for explanatory pluralism.

20 I presented this paper at Arizona State University, the Leeds-ACU Normativity Workshop, and the 2023 Workshop in Normative Ethics in Tucson. Many thanks to these audiences, as well as Sam Baron, David McElhoes, Alex Stamatiadis-Bréhier, and referees for this volume, for their very helpful feedback.

References

Bader, Ralf (2017), "The Grounding Argument Against Non-Reductive Moral Realism", *Oxford Studies in Metaethics* 12: 106-134.

Baker, Derek (2021), "If You're Quasi-Explaining, You're Quasi-Losing", *Oxford Studies in Metaethics* 16: 54-79.

Baron, Sam and Norton, James (2021), "Metaphysical Explanation: The Kitcher Picture", *Erkenntnis* 86: 187-207.

Berker, Selim (2018), "The Unity of Grounding", *Mind* 127: 729–777.

Berker, Selim (2019), "The Explanatory Ambitions of Moral Principles", *Noûs* 53, 904-936.

Brueckner, Anthony L. and Fischer, John Martin (1986), "Why Is Death Bad?", *Philosophical Studies* 50: 213-221.

Chappell, Richard Yetter (2021), "The Right Wrong-Makers", *Philosophy and Phenomenological Research* 103: 426-440.

Cullity, Garrett (2018), *Concern, Respect, and Cooperation* (Oxford University Press).

Duguid, Callum (2021), "Lawful Humean Explanations Are Not Circular", *Synthese* 199: 6039-6059.

Enoch, David (2019), "How Principles Ground", *Oxford Studies in Metaethics* 14: 1-22.

Fabre, Cécile (2003), "Justice and the Compulsory Taking of Live Body Parts", *Utilitas* 15: 127-150.

Faulkner, Paul (2007), "What Is Wrong with Lying?", *Philosophy and Phenomenological Research* 75: 535-557.

Fine, Kit (2012), "Guide to Ground", in *Metaphysical Grounding*, ed. F. Correia and B. Schnieder, 37-80 *(*Cambridge University Press).

Finlay, Stephen (2019), "Defining Normativity", in *Dimensions of Normativity*, ed. D Plunkett, S. Shapiro, and K. Toh, 187-219 (Oxford University Press).

Fogal, Daniel and Risberg, Olle (2020), "The Metaphysics of Moral Explanations", *Oxford Studies in Metaethics* 15: 170-194.

Gijsbers, Victor (2007). "Why Unification is Neither Necessary nor Sufficient for Explanation", *Philosophy of Science* 74: 481–500.

Hempel, Carl G. and Oppenheim, Paul (1948), "Studies in the Logic of Explanation", *Philosophy of Science* 15: 135–175.

Kamm, F. M. (1993), *Morality, Mortality Volume I: Death and Whom to Save From It* (Oxford University Press).

Khalifa, Kareem, Doble, Gabriel, and Millson, Jared (2020), "Counterfactuals and Explanatory Pluralism", *British Journal for the Philosophy of Science* 71: 1439-1460.

Kitcher, Philip (1989), “Explanatory Unification and the Causal Structure of the World”, in *Scientific Explanation*, ed. P. Kitcher and W. Salmon, 410–505 (University of Minnesota Press).

Korsgaard, Christine M. (2007), “What’s Wrong with Lying?”, in *Philosophical Inquiry: Classic and Contemporary Readings*, ed. J. E. Adler and C. Z. Elgin, 577-587 (Hackett).

Kovacs, David Mark (2020), “Metaphysically Explanatory Unification”, *Philosophical Studies* 177: 1659-1683.

Lange, Marc (2016), *Because Without Cause: Non-causal Explanations in Science and Mathematics* (Oxford University Press).

Lange, Marc (2023), “A New Circularity in Explanations by Humean Laws of Nature”, *Philosophical Studies* 180: 1001-1016.

Lewis, David (1986), “Causal Explanation”, in his *Philosophical Papers: Volume II*, 214-40 (Oxford University Press).

Mantzavinos, Chrysostomos (2016), *Explanatory Pluralism* (Cambridge University Press).

Mayer, Robert (2007), “What’s Wrong with Exploitation?”, *Journal of Applied Philosophy* 24: 137-150.

Pincock, Christopher (2018), “Accommodating Explanatory Pluralism”, in *Explanation Beyond Causation*, ed. A. Reutlinger and J. Saatsi, 39–56 (Oxford University Press).

Potochnik, Angela (2015), “Causal Patterns and Adequate Explanations”, *Philosophical Studies* 172: 1163-1182.

Potochnik, Angela (2020), “What Constitutes an Explanation in Biology”, in *Philosophy of Science for Biologists*, ed. K. Kampurakis and T. Uller, 21-35 (Cambridge University Press).

Rice, Collin and Rohwer, Yasha (2021), “How to Reconcile a Unified Account of Explanation with Explanatory Diversity”, *Foundations of Science* 26: 1025-1047.

Robeyns, Ingrid (2022), “Why Limitarianism?” *Journal of Political Philosophy* 30: 249-270.

Roemer, John (1996), *Egalitarian Perspectives* (Cambridge University Press).

Rosen, Gideon (2017), “Ground by Law”, *Philosophical Issues* 27: 279–301.

Ross, W. David (1930), *The Right and the Good* (Clarendon Press).

Ruben, David-Hillel (2012), *Explaining Explanation*, updated and expanded second edition (Paradigm Publishers).

Sachs, Benjamin (2018), *Explaining Right and Wrong* (Routledge).

Schurz, Gerhard (1999), “Explanation as Unification”, *Synthese*, 120: 95–114.

Sinnott-Armstrong, Walter and Miller, Franklin G. (2013), “What Makes Killing Wrong?”, *Journal of Medical Ethics* 39: 3-7.

Skiles, Alexander (2015), “Against Grounding Necessitarianism”, *Erkenntnis* 80: 717-751.

Stamatiadis-Bréhier, Alexios (forthcoming), “Grounding Functionalism and Explanatory Unificationism”, *Journal of the American Philosophical Association.* URL = <https://doi.org/10.1017/apa.2022.29>.

Stamatiadis-Bréhier, Alexios (ms), “How Moral Generalizations Explain”, unpublished.

Sterri, Aksel Braanen (2021), “Why States Should Buy Kidneys”, *Journal of Applied Philosophy* 38: 844-856.

Strevens, Michael (2004), “The Causal and Unification Approaches to Explanation Unified—Causally”, *Noûs* 38: 154-176.

Strevens (2008), *Depth: An Account of Scientific Explanation* (Harvard University Press).

Stroud, Sarah (2024). “How Do Personal Relationships Make a Moral Difference?”, *Oxford Studies in Normative Ethics*, this volume.

Valdman, Mikhail (2009), “A Theory of Wrongful Exploitation”, *Philosophers’ Imprint* 9: 1–14.

Väyrynen, Pekka (2021a), “Normative Explanation and Justification”, *Noûs* 55: 3-22.

Väyrynen, Pekka (2021b), “Normative Explanation Unchained”, *Philosophy and Phenomenological Research* 103: 278-297.

Vrousalis, Nicholas (2013), “Exploitation, Vulnerability, and Social Domination”, *Philosophy and Public Affairs* 41; 131-157.

Walden, Kenneth (2016), “The Relativity of Ethical Explanation”, *Oxford Studies in Normative Ethics* 6: 189-213.

Wilson, James Lindley (2022), “Making the All-Affected Principle Safe for Democracy”, *Philosophy and Public Affairs* 50: 169-201.

Wood, Allen (1995), “Exploitation”, *Social Philosophy and Policy* 12: 136-158.

Woodward, James (2003), *Making Things Happen* (Oxford University Press).