



# Moral Generalizations and Moral Explanatory Pluralism

Alexios Stamatiadis-Bréhier<sup>1,2</sup>

Received: 22 February 2024 / Accepted: 9 September 2024  
© The Author(s) 2024

## Abstract

I argue that moral principles, construed as moral generalizations, can be genuinely explanatory. Specifically, I present and respond to a challenge according to which moral generalizations are explanatorily redundant. In response, I present and defend an explanatory dimension of moral generalizations that is based on the idea of unification. I do so in the context of motivating a realist-friendly moral explanatory pluralism (i.e., the view that there can be many, equally legitimate, explanations of moral facts). Finally, I appeal to the same theoretical resources to tackle an objection from explanatory circularity.

## 1 Introduction

In this paper, I argue that moral principles, construed as moral generalizations, can be genuinely explanatory. Moral generalizations are lists of moral phenomena in the sense that they are nothing over and above the instances that fall within their scope. The core principle of consequentialism, so construed, would consist in all and only those instances of wrongness (possible and actual) that hold in virtue of their non-moral constituents (which presumably contribute towards some specified consequence). In this sense, moral principles *just are* the sum of their particular instances.

This is a *non-substantivist* view concerning the nature of moral principles: moral principles reduce to (or are fully grounded by) the instances that comprise them.<sup>1</sup>

<sup>1</sup> There are two technical issues I want to bracket since they are not directly relevant to neither the goal nor the dialectic of this paper. First, I loosely oscillate between grounding and reduction-talk in my formulation of moral generalizations. The link between reduction and grounding is a thorny issue (for discussion, see Dorsey (2016) and Rosen (2017b)) but the main point I am making here is clear enough: as per the non-substantivist view, moral principles are not distinct entities that *govern* their instances. Secondly, it is likely that moral generalizations reduce to (or are fully grounded by) their instances *plus* a totality fact (and/or a “that’s it” clause) (for similar discussion on the notion of “minimalist” grounding laws, see Giannotti (2022); for an unorthodox view that rejects that generalizations are grounded in their instances, see Kovacs (2020a, 2021) and Marshall (2015)).

✉ Alexios Stamatiadis-Bréhier  
alexios.brehier@gmail.com

<sup>1</sup> Azrieli International Postdoctoral Fellow, Department of Philosophy, Tel Aviv University, Tel Aviv, Israel

<sup>2</sup> Postdoctoral Fellow, Department of History and Philosophy of Science, National Kapodistrian University of Athens, Athens, Greece

Compare with *substantivism* about moral principles. Such views take moral principles to be something *over and above* their instances. In fact, they take moral principles to *govern* these instances in the same way non-Humean views of scientific laws are taken to govern particular causal relations. For example, Rosen (2017a, b) has suggested that moral principles involve relations between *universals*, inspired by the so-called Dretske-Armstrong-Tooley account of scientific laws (for discussion, see Bhogal (2017)).

If, as per the non-substantivist view, a moral principle is just a sum of its actual and possible instances, how can it be *explanatory* of those instances? It could be argued that moral generalizations are, by their nature, explanatorily idle.<sup>2</sup> Based on plausible assumptions about the form of moral principles, moral generalizations specify the full grounds of particular moral facts (see Section 2). But then it seems that particular moral facts are fully explained in terms of their full grounds, thus rendering the relevant moral principle explanatorily obsolete.

The view that moral principles are generalizations in the non-substantivist sense might seem implausible to some. Still, it is worth defending for at least two reasons. First, it would follow from a broadly neo-Humean metaphysical perspective. A neo-Humean account of moral principles would be the view that moral principles are nothing over and above particular entities figuring in the Humean mosaic (roughly, a base of highly specific, non-modal, categorical properties) (Lewis, 1994). Neo-Humeanism is a strong contender in the literature on scientific laws and a powerful metaphysical framework overall. Its successes can be found in multiple areas in philosophy, ranging from metaphysics of naturalness (Lewis, 1983) to moral metasemantics (Dunaway and McPherson, 2016). At the very least, neo-Humeanism about moral principles should be understood as a live conceptual possibility that is worth exploring.<sup>3</sup>

Secondly, the view that moral principles just are the sum of their instances is superior to rival views in terms of ontological parsimony and philosophical parochiality. Views that understand moral principles as something *distinct* from their instances need to further specify the nature of that separate entity and countenance the relevant theoretical cost that comes with its adoption. For example, consider a view according to which moral principles are metaphysically *primitive*, analogous to the nomic primitivist view concerning the laws of nature (e.g., Maudlin (2007)). Such moral principles would be fundamentally distinct from the entities that fall within their scope: these principles *govern* their instances and, in this sense, are metaphysically *separate* from them (recall Rosen's view). Other similar views appeal to essences, relations between universals, or other metaphysical posits (for discussions, see Fogal and Risberg (2020)).

Why would one want to defend the genuine *explanatoriness* of moral generalizations? There are two main reasons. First, there is a consensus in first-order normative ethics that moral principles are explanatorily important.<sup>4</sup> Explanations of particular moral facts somehow involve an appeal to some general moral fact. A particular act

<sup>2</sup> See Berker (2018) and Robinson (2007: sec. 3; 2011: 295–6) (although see fn. 10).

<sup>3</sup> Robinson (2007; 2011) and Fogal and Risberg (2020) briefly consider this view (see also Berker (2018)).

<sup>4</sup> For references, see Robinson (2011: 290; 294 fn. 7), Walden (2016), and Fogal and Risberg (2020).

is either good or bad partially because a specific moral principle holds. For example, a particular killing is morally wrong partially because such an act decreases utility (*ceteris paribus*) and, in general, moral goodness consists in utility. If moral generalizations are to be considered moral *principles*, then they should be able to accommodate this datum.<sup>5</sup>

Secondly, *scientific* laws are also routinely used to explain particular scientific facts. Many theories of scientific explanation take this to be uncontroversial (e.g., Woodward and Hitchcock (2003)). Insofar as the analogy between scientific and moral laws is warranted, we should expect moral principles to work in a similar way. Of course, that analogy can be resisted. But the people who typically accept that moral principles are like scientific laws are the same people who would be interested in defending the explanatoriness of moral principles. For example, realist moral naturalists construe the set of moral facts as a proper subset of the set of natural facts. In this sense, it is natural for them to take moral principles to be akin to scientific laws (this view is explicitly endorsed by Stamatiadis-Bréhier, 2022). The analogy is also typically granted by realist non-naturalists. After all, there are important precedents in appealing to moral principles to explain moral supervenience in the same way laws concerning mentality can be used to explain the supervenience of the mental on the physical (see McPherson (2012)).

In the next section, I present the *redundancy challenge* against the explanatoriness of moral generalizations (Section 2). In response, I suggest that moral generalizations can explain via subsumption (Section 3). This view is further vindicated by showing that it is a part of larger theory on moral explanation: moral explanatory pluralism (Sections 4 and 5). Finally, in Section 6, I anticipate an objection to the view that moral generalizations explain via subsumption: what I call the *circularity challenge* (Rosen, 2017a, b). Roughly, even if somehow moral generalizations can explain particular moral facts, there is a revenge problem lurking in the background. It is plausible that moral generalizations are explained *by* their instances (given that they reduce to, or are fully grounded by, them). So, it seems that, on

<sup>5</sup> Note that, in a non-substantivist setting, expressions of this sort should be read felicitously. Saying that “goodness consists in pleasure,” for example, should be understood as saying that the sum of instances of goodness modally covary, and are metaphysically determined by (see Section 2), equivalent instances of pleasure (where such instances include both actual and possible ones). This illustrates that talk of “goodness” does not commit the non-substantivist to a *universal* of goodness (ditto for pleasure). Perhaps at this point, one could object that instances of pleasure must have something in common and a prime candidate for making sense relations of objective similarity between such instances is a universal. But this seems to be in tension with the non-substantivist view. Thankfully, this result can be avoided. The non-substantivist can appeal to a *nominalist* construal of properties to accommodate objective similarity between instances without reifying properties in the form of universals. For example, there are plausible accounts in the literature about construing moral properties in the context of trope theory (e.g., see Ridge (2007), Suikkanen (2024a, b), and Sinclair (2024) for recent discussion) (other options involve adopting some kind of resemblance nominalism as per Rodriguez-Pereyra (2006)). At any rate, it seems that this is not a problem for non-substantivism per se, but for any view that wishes to account for objective similarity between moral instances without adopting universals (to be clear, even adopting an ontology of tropes would be a significant theoretical cost, although in my view this would still be compatible with the minimalist view as stated). Also, given the focus of my paper, largely sidelining issues which are not directly relevant to the *explanatoriness* of moral generalizations is dialectically acceptable on my end. Many thanks to an anonymous referee for some very helpful discussion.

pain of circularity, we cannot appeal to those generalizations to explain those very instances. In response, I use the explanatory pluralist machinery once more to dispel the circularity challenge.

## 2 The Redundancy Challenge

Consider the following view:

**SIMPLE:**  $M_1$  (a particular moral fact) is partially explained by  $N_1$  (a particular non-moral fact) and partially explained by  $G$  (a moral generalization).

According to what I will call the *redundancy challenge*, a moral generalization ( $G$ ) is explanatorily redundant towards a moral instance falling within its scope ( $M_1$ ).

A moral principle states a connection between a general moral fact and a general non-moral fact (i.e., a fact whose particular instances exemplify it). For example, take once again the principle according to which moral goodness consists in pleasure. This relation should not be understood in purely modal terms: it is not simply the case that moral goodness ( $M$ ) modally covaries with pleasure ( $N$ ). Rather,  $M$  obtains fully *because* of  $N$ . In this sense,  $M$  is fully *explained* by  $N$  (for this point, see Berker (2018)).

Now take particular moral facts. For example, if the wrongness of killing holds fully in virtue of the descriptive features of such acts, then the wrongness of *particular* killings also holds fully in virtue of *their* descriptive features. In this sense, the explanation of the wrongness of a particular killing in terms of its descriptive features corresponds to a full explanation of the form “The wrongness of this particular killing is fully explained by its descriptive features.” But if moral instances are fully explained in terms of their corresponding descriptive features, then  $G$  seems to become explanatorily obsolete. For this reason, according to the redundancy challenge, **SIMPLE** is false.

What is the notion of redundancy at play here? Here is what I take to be a plausible suggestion: something is not redundant when it is *needed*. So, plausibly, a ground is redundant if it not *needed* for the target fact to obtain. A plausible way to read this claim is to say that a ground  $P$  is redundant for a target fact  $Q$  when  $Q$  obtains, in some sense, *regardless* of whether  $P$  obtains. What is the notion of ground at play here? Even under minimal assumptions about grounding, it is commonly accepted that grounds contribute towards making their target facts obtain (e.g., Litland (2013)). In this sense, a full ground is *sufficient* for its target fact.

There are many different ways in which a ground can have a contributory role. Let us say that, on the one end of the spectrum, a ground can be *crucially contributory* in the sense that its contribution is particularly important for the fact it grounds. For example, one could claim that many factors contributed towards some organisms having phenotype  $P$ , but the most important factor for having  $P$  is having gene  $Q$ . On the other end, a ground can be *non-crucially contributory* in the sense that it is not *needed* for its target fact to obtain; many other facts could have figured in its place without altering the result. For example, my injury was *crucially* caused by the fact that I was hit by an assailant, but it was *non-crucially* caused by the fact that I was hit by a *blonde* assailant (see Woods (2018) on vacuous grounding). Or, there might be

cases where a factor is not needed in the sense that it is *trivial*: the fact that I exist can be non-crucially explained by the fact that the world exists, and I am a component of that world.

A plausible way to read the redundancy challenge is by adopting the distinction between crucial and non-crucial grounds. Redundant grounds are non-crucial grounds and non-redundant grounds are crucial grounds.<sup>6</sup> The objector is correct in saying that, in the context of **SIMPLE**, G is less crucial than  $N_1$ , towards  $M_1$ .  $N_1$  is more important than G in the following sense: even though both facts, strictly speaking, fully contribute towards making  $M_1$  the case, G has many *idle parts*.

To see this, consider G's form. As mentioned, G fully reduces to the list of every instance falling within G. A moral generalization about the wrongness of killing is simply a summary of every particular killing (actual *and* possible) alongside the non-moral features that brought them about (and relevant instances of grounding relations).<sup>7</sup> In that respect, G involves  $N_1$  which, by itself, is a full ground for  $M_1$ , while also involving the grounds of *other* instances. Even though G, on its own, fully contributes towards  $M_1$ , it does so in a non-crucial way: other facts having  $N_1$  as their part would also bring about  $M_1$ . In that respect, there is nothing particularly *special* about G.

### 3 Subsumption Explanations

The explanatory intuition that motivates the redundancy challenge seems to be the following:

**NON-IDLENESS:** The set of explanantia should involve the non-idle entities that brought the explanandum fact into existence.

And, to be clear, there is certainly something true about **NON-IDLENESS**. So-called *mechanistic* explanations focus on the fine-grained details that lead from a mechanism's

<sup>6</sup> There are many possible fine-grained ways to understand crucial grounds which I am not going to rehearse here. I take the distinction between crucial and non-crucial grounding to be intuitive and independently plausible. My argument relies only on a general characterization of this distinction (for a plausible development of this idea, see Strevens's (2008) *kairctic* account). More fleshed out accounts have been offered by Krämer and Roski (2017), Woods (2018), and Makin (2019). For general discussion on the varieties of "non-causally making it the case," see Wygoda Cohen (2020: sec. 2).

<sup>7</sup> For the purposes of this paper, I take the inclusion of "because" and "grounding" locutions in the formulation of moral principles to be a plausible and (dare I say) non-parochial assumption. It is widely accepted in the literature (e.g., Fogal and Risberg (2020)), including those having advanced similar arguments to the redundancy challenge (e.g., Berker (2018)) (although see fn. 10). And this is unsurprising given the multiple philosophical benefits of hyperintensional formulations of moral principles (e.g., addressing Euthyphro-style questions). It is also compatible with the non-substantivist framework, which takes a given moral principle *to just be* the sum of its instances (bracketing fn. 1). In this sense, an instance of the form " $Mx_1$  iff, and because,  $Nx_1$ " partially grounds a moral principle of the form " $Mx$  iff, and because,  $Nx$ ." The presence of a determination/grounding relation signalled by the "because" locution is compatible with non-sustantivism insofar it concerns the *instances* of such relations figuring in the scope of a given moral principle. Granted, the neo-Humean theory of moral principles (a particular *version* of non-substantivism) does have some work to do concerning the accommodation of the grounding relation (although see Sider (2020)) but this is something that I cannot examine in this paper (which, recall, concerns specifically the claim about the explanatoriness of moral generalizations). Many thanks to an anonymous referee for related discussion.

set-up to its termination conditions (Machamer et al., 2000). For example, the mechanistic explanation for a given event of protein-synthesis involves a pre-initiation sequence (roughly, an mRNA and a tRNA coming together) and a termination phase (roughly, the mRNA degrades so the relevant nucleotides can be used elsewhere for other processes). It is for this reason that some accounts of mechanistic explanation do not include laws in the relevant set of explanantia (cf. Cartwright et al., 2020). Such explanations seem to get their explanatory force from the fact that they involve the components, and only those components, that *actively* brought about the relevant phenomenon. The requirement stated by **NON-IDLENESS** seems to capture this insight.

Still, I want to challenge the inference from the fact that G non-crucially grounds  $M_1$  to the claim that G cannot be *explanatory* towards  $M_1$ . Specifically, even though “G explains  $M_1$ ” is not an explanation in terms of **NON-IDLENESS**, it does not follow that it cannot be an explanation of another kind. Indeed, as I will argue, it is independently plausible that **NON-IDLENESS** is only *one* of the (potentially many) constraints under which a moral explanation can operate.

Consider what I will call *subsumption* or *unification-based* explanations. Such explanations explain a particular phenomenon by showing that it is a part of a larger whole.<sup>8</sup> For example, it is plausible that one could explain the wrongness of a particular killing in virtue of the fact that all killings have such-and-such features, thus showing that the explanandum fact is a part of a larger pattern (see Kitcher (1989) for a classic development of the unificationist framework) (for recent discussion on ‘non-causal’ unificationism see Kovacs, 2020b; Stamatiadis-Bréhier, 2023). If utilitarianism is true, then the fact that the moral goodness of an act  $\varphi$  holds in virtue of the fact that  $\varphi$  contributes towards utility maximization is not an isolated phenomenon: *other* instances qualitatively similar to  $\varphi$  would *also* be instances of moral goodness. These explanations aim for unification: they get their explanatory force by showing how the relevant phenomena fit together.

**UNIFICATION:** The set of explanantia should involve a general entity under which the explanandum fact falls. Such explanations involve the explanation of a particular phenomenon by showing how that phenomenon is a part of a larger pattern.

In this sense, **SIMPLE** seems to be compatible with **UNIFICATION**.<sup>9</sup>

One might worry that, as a response to the redundancy challenge, the appeal to **UNIFICATION** misses the mark. Presumably, the redundancy challenge is framed in terms of a specific sense of explanation: explanation that is backed by crucial grounding. So the fact that there is *another* kind of explanation which accommodates **SIMPLE** is neither here nor there.

I should stress, however, that I do not take issue with the view that G is explanatorily redundant *under* **NON-IDLENESS**. Rather, I take issue with the thesis that

<sup>8</sup> A subtle point here concerns the exact profile of the explanandum. If Q is subsumed by P, then Q is explained by P. Alternatively, one could also explain P’s *place* in Q in terms of Q. What is at stake here is the former kind explanation, not the later.

<sup>9</sup> Baker’s (2021) recent suggestion builds on a similar notion of subsumption. But his account is significantly different from the one presented in this paper since he appeals to a *pragmatic* theory of explanation in order to provide an expressivist-friendly construal of the “because” locution (cf. Baker 2021: 56).

**NON-IDLENESS** is a *requirement* or, less strictly, a more *important* notion of explanation compared to **UNIFICATION**.<sup>10</sup> In other words, it is unclear to me in what sense one should care more about explanations that fall under **NON-IDLENESS** compared to explanations that fall under **UNIFICATION**. And if there is no such reason, then the severity of the redundancy challenge is greatly diminished.

Perhaps there *is* something problematic with **UNIFICATION** which makes explanations falling under its scope somehow inferior compared to explanations that fall under **NON-IDLENESS**. First, there might be some suspicion concerning the *metaphysics* of subsumption explanations. It is usually suggested that explanations must be *objective* as per the *ontic conception* of explanation. According to this conception, genuine explanations are objective by involving entities that exist in a stance-independent way (Craver, 2009). Consider a type of explanation that does not meet that requirement. Take, for example, a set of explanantia which involve abstractions in the form of idealizations (e.g., Strevens (2008: ch. 8)). Such abstractions might be epistemically useful but, according to the ontic conception, they are not *genuinely* explanatory since they are not objective features of the world.

But even if fictional explanantia are objectionable, moral generalizations like G are not like fictions. As per non-substantivism, G is the sum of its instances, and those very instances are fully objective, non-fictional, entities. According to the ontic conception, explanations need to involve nothing but the objective features of the world. But there are *many* ways in which this requirement can be met. One way is by appealing to an explanans with idle parts like G. Another is to only appeal to those components of G that are crucial grounds for the relevant explanandum fact. If the ontic conception is simply a constraint concerning the stance-independent existence of the explanantia facts, then **NON-IDLENESS** is not the *only* way that such a constraint can be accommodated.

On a different interpretation of the ontic conception, it could be argued that explanations should not involve any *epistemic* components (for discussion, see Craver (2009)). That is, explanations should be 'backed' by *purely* stance-independent entities (Stamatiadis-Bréhier, 2021). But I find such a version of the ontic conception to be implausibly stringent. If there is a ban on *any kind* of epistemic, pragmatic, or otherwise relativization components into the success conditions of an explanation, then the explanations that *do* satisfy the ontic conception will be akin to something like Railton's "Ideal Explanatory Text": a story involving *every* single, minute or not, contributing factor that brought about the relevant explanandum-phenomenon (Railton, 1980).

Secondly, it seems to me that *even NON-IDLENESS* fails under such an understanding of the ontic conception. Let it be a requirement on explanations that they should be backed only by entities involving non-idle parts. To my mind, this is *also* an epistemic way of constraining the total set of contributing factors that brought

<sup>10</sup> This is why my paper is *not* supposed to be understood as a reply to Berker (2018) (cf. Stamatiadis-Bréhier (unpublished manuscript)). Berker's challenge is similar to the redundancy challenge as stated, but it is presented in a much narrower sense. Specifically, Berker argues against the view that moral generalizations ground their instances in terms of what I've labelled "crucial grounding" (a conclusion which, as noted, I grant). Presumably this would still allow Berker to accept that moral generalizations ground (and thus explain) their instances in a non-crucial way (e.g., as per **UNIFICATION**).



about a certain phenomenon (without this suggesting that the relevant explanation is somehow rendered non-objective or genuine).

So, to conclude, it seems that there is no plausible way of claiming that explanations that fall under **NON-IDLENESS** are, in principle, in any way more important or legitimate than explanation that fall under **UNIFICATION**. Of course, the application of **UNIFICATION** to *moral* facts is something that needs to be motivated separately (and I do just that in the following sections).

## 4 Towards Moral Explanatory Pluralism

One can have it both ways by adopting both **NON-IDLENESS** and **UNIFICATION**. This can be done by adopting a pluralistic theory of moral explanation which individuates the relevance constraints of every explanation in terms of what is decided by the appropriate research project. In this sense, an explanans is redundant only in the context of an explanatory project:

**PLURALISM:** P explains Q, iff, P is an appropriate explanans for Q given the aims and goals of the research project in which the explanation “P explains Q” figures.

In this paper, I will not attempt to present a complete theory behind **PLURALISM**. Rather, I take **PLURALISM** to be intuitively plausible and, at the very least, an interesting philosophical view concerning the phenomenon of normative explanation that deserves serious consideration (for recent discussion see Väyrynen forthcoming). Still, to put some flesh on the bones of **PLURALISM**, I will make the following two notes.

First, the suggestion that a research project determines the appropriateness of a putative explanans requires some elucidation. The idea is that for any given phenomenon P, there is a multitude of factors (causal and non-causal) which contribute towards making P the case (recall Railton’s Ideal Explanatory Text). **PLURALISM** suggests that there are many different ways of *selecting* which set of facts are *appropriate* for explaining P, since different criteria of selection are useful for different *purposes*. To illustrate, consider how Potochnik (2015) spells out this sort of idea in the context of scientific explanation:

The research program in which an explanation is formulated accounts for the [relevant] contextual influence. An explanation of an event or phenomenon always occurs in the context of some research agenda. This narrows the scope of investigation to a segment of causal history, brings certain types of causal relationships to the fore, and involves certain methods of investigation and representation. (2015: 12; modified)

For example, the explanation of Harris sparrows exhibiting specific color variation within a single flock can have many different explanations depending on the contributing factors that one chooses to focus on (as indicated by the relevant research project). Evolutionary game theory would focus on the *status* of certain sparrows in terms of the available food resources and its influence towards their color variation, whereas a



different biological program would explain that same color variation by focusing on *genetic* factors (see Potochnik (2015: 12–3)). Of course, whether a research project is warranted in focusing on some specified set of contributing factors can be a matter of controversy. But under the plausible assumption that there are *multiple* legitimate research programs about P, it follows under **PLURALISM** that there are many different equally legitimate *explanations* about P.

Secondly, we should resist the idea that the various relativization components that **PLURALISM** can introduce (given the different ways the appropriateness of an explanation can be construed) would somehow compromise the objectivity of these explanations. On the contrary, **PLURALISM** is well-fit to be a *realist* theory of explanation. The fact that a research group has some set of aims does not necessary entail that these aims are going to be accommodated. On a realist picture, these aims will be accommodated only if the entities specified by these aims actually exist. To echo Boyd (1999: 89), explanation is *bicameral*: both the mind-independent world and our epistemic practices need to come together to make an explanation successful. Thankfully, the same holds under **PLURALISM**: it is just that there are *many* ways in which our epistemic practices can contribute to the success conditions of an explanation.

**NON-IDLENESS** can be incorporated under the moral explanatory pluralist picture I have just sketched. The proponent of the redundancy challenge is indeed correct in claiming that moral generalizations are explanatorily redundant if one assumes **NON-IDLENESS**: the explanandum fact (a particular moral fact) is fully explained (in the relevant sense), without residue, in virtue of those particular facts that crucially contributed towards making that fact the case.

What research aim could **NON-IDLENESS** plausibly serve? One plausible possibility concerns the fact that **NON-IDLENESS** explanations demonstrate a highly specific level of *grain*. Their function is to showcase in great detail what brought about a given explanandum fact: the moral wrongness of a particular act was non-causally brought about by such-and-such particular facts. What explains the fact that Richard Nixon acted immorally? Answer: the fact that he was a central component of a complicated causal sequence which involved individuals such as John Mitchell, John Dean, and Gordon Liddy (who set up “Operation Gemstone”), and events such as the secret “hush fund” as part of the Committee to Re-elect the President, which eventually resulted in what is now known as the “Watergate scandal.”

Analogously, an event involving the synthesis of a protein is explained by those fine-grained components that contributed towards it. Still, as already argued, showing that some explanations are plausibly constrained by **NON-IDLENESS** does not entail that *every* explanation should be as such. In the next section, I will further defend the claim that **UNIFICATION** explanations can be incorporated under **PLURALISM** and thus be genuine explanations.

## 5 UNIFICATION Again

What about **UNIFICATION**? I submit that one plausible aim which **UNIFICATION**-based explanations can serve is that of *prediction*. Note that even though explanation and prediction are different enterprises, some explanations are more

useful than others in terms of facilitating some predictive goals. For example, a moral explanation falling under **NON-IDLENESS** does not provide enough information that would allow us to predict what would happen if similar circumstances were to obtain: these explanations only state which fine-grained contributing factors *actually* brought about a particular moral fact. On the other hand, knowing that, *in general*, wrongness is brought about by descriptive facts that adhere to the Categorical Imperative is much more useful in that regard. Situating a particular grounding fact (e.g., “[N<sub>1</sub> grounds M<sub>1</sub>]”) into a larger pattern allows us to predict what would happen if similar but uninstantiated instances of that pattern were to obtain.

The fact that moral principles have the function to predict seems to be more connected to their function as guides rather than standards. But this need not be the case. Moral principles can be understood as standards which explain, and, *in turn*, one can understand prediction as one of the potentially *many* functions that explanation can exhibit. So, one can appeal to moral principles as standards to predict, but only *derivatively*, in virtue of explanation *itself* having a predictive function (in some of its instances).

Explanation and prediction are importantly distinct enterprises since they can come apart. Still, there is a substantive, non-trivial, link between prediction and certain forms of explanation. Woodward (2003), for example, develops an interventionist theory of explanation according to which if P explains Q then a counterfactual of the form “Had one intervened upon the P-variable, this would result in a change in the Q-variable” is true. The idea behind this account is that a primary motivation behind our explanatory practices is our desire to *control*, and thus predict, various aspects of the phenomena around us (Woodward, 2003: 12).

I take it that the above is plausibly true about explanations of the form “P<sub>1</sub> explains Q<sub>1</sub>” where P<sub>1</sub> and Q<sub>1</sub> are singular tokens. But if this is the case, then the same should be the case for facts involving a *general* explanans and a particular explanandum falling under the former. As already mentioned, if I know that P<sub>1</sub> brings about Q<sub>1</sub> and that it is *generally* the case that entities *like* P<sub>1</sub> also bring about entities like Q<sub>1</sub>, then my ability to control such phenomena is significantly enhanced. So, even though this kind of predictive dimension of explanation is not *unique* to **UNIFICATION**, the *scope* of explanations that proceed by unifying allows for a wide range of predictions to become available.

More concretely, **UNIFICATION** explanations can be plausibly identified in the context of existing metaethical theories specifically. Boyd (1988: 329) takes moral goodness to be grounded in a cluster of important human needs that is sustained by a homeostatic mechanism. Crucially, according to Boyd, this mechanism has a hybrid nature: its proper functioning depends on factors that are both *internal* and *external* to the nature of human needs (Boyd, 1988: 330–1). So, it seems plausible that there will be important counterfactually robust generalizations about the conditions under which these needs can be met. These generalizations will concern important details about both the internal characteristic of human needs, as well as the societal and theoretical resources that are necessary for their satisfaction. A **UNIFICATION**-based explanation would explain a particular instance of moral goodness in terms of the fact that it is grounded in terms of natural facts that fall within these internal/external satisfaction conditions.

Another goal that **UNIFICATION** can serve concerns our *understanding* of moral goodness. Consider monist attempts to bring together various moral principles under the same notion of goodness. Mill famously subsumes lower and higher level types of pleasure under his utilitarian principle (Brink, 2013). Ross initially lists seven *prima facie* duties, but later suggests that some of them (beneficence, self-improvement, and justice) can be understood as a *single* duty concerning intrinsic goodness (Ross, 1939). These examples are controversial, but they work as illustrations. In both cases, regardless of the relevant consequences on first-order normative discourse, the subsumption of less fundamental principles under more fundamental ones significantly increases one's understanding of goodness. After all, we are not only interested in what we should do (morally speaking) but also in the nature of goodness itself. **UNIFICATION** works towards that direction.

A worry about the above illustrations is that they involve the explanation of a normative entity in terms of *another* normative entity (e.g., the explanation of a moral principle in terms of *another*, more general, moral principle). But this should not alarm us since the same type of explanation can be situated in terms of an instance and the general normative entity under which it is subsumed. For example, knowing the fact that a particular instance of moral wrongness falls under the principle of (say) utility increases our understanding of *that very principle*. We know more about the principle of utility and its implications once we have a broader understanding of the instances that fall within its scope. To see this, assume that a relatively controversial instance of moral wrongness (e.g., our duty towards people of the distant future) falls under the principle of utility. In that case, identifying that instance in the context of the principle of utility significantly contributes towards our understanding of that principle and, in turn, of the nature of moral goodness itself.

One final objection. An anonymous reviewer raises some worries about **UNIFICATION** and whether it satisfies the theoretical virtue of prediction in a *distinctive* way. They argue that a moral explanation satisfying **NON-IDLENESS** would *also* have a predictive function. Consider the following explanation:

**CASE:** Tony's killing Anne was wrong (M), and the wrongness of his action is fully grounded in the fact that he had a clear intention of killing an innocent person (N)

M and N are, respectively, moral and natural properties. At this point, the objector asks: Wouldn't the truth of **CASE** be enough to predict that relevantly similar actions, such as Katia's killing someone with a clear intention to kill an innocent person make her action wrong? If so, it seems that moral principles of the form "killing someone with a clear intention of killing an innocent person is wrong" is indeed redundant, *even* in the context of subsumption or unification-based explanations. Why? Because prediction can occur even in the *absence* of moral principles (qua generalizations).

In response, I argue that **CASE**, on its own, is indeed unable to explain how other qualitatively relevantly similar actions are also morally wrong (assuming that they are as such). The reason for this is the possibility of *particularism*. Roughly, due to the possibility of *reasons holism*, it could be that the instantiation of a set of natural properties in some context C does not ground the same moral fact in context C\*. In this sense, it

could be that Tony's killing is morally wrong, even though Katia's killing is morally right (perhaps due to the lack of certain enablers) (see Väyrynen (2023)).

Effectively what I am suggesting is that even though it might *look as if* **CASE** is an explanation with a predictive function, what is actually doing the work is the underlying assumption of moral generalism. And quite clearly, since moral generalism *just is* the view that there are moral principles at play, it is thanks to *their* scope that **CASE** can be used for prediction (assuming that it can be).

It could be objected that the possibility of particularism simply illustrates the need for moral *principles*, not moral *generalizations* (at least as far as explanatory unification is concerned). The worry here is that perhaps only *substantive* moral principles can serve our predictive goals, given that only substantive accounts can *explain* the counterfactual robustness of such principles. For example, it is commonly asserted among substantivist circles that moral principles owe their *generality* and *necessity* to underlying facts about essences, or perhaps primitivist nomic facts (e.g., see Haderlie and Litland forthcoming). To compare, moral generalizations, as per the non-substantivist account I have been advancing so far, are not underwritten by such extravagant metaphysical posits. Instead, moral principles are simply taken to be grounded in their instances, and that's it.

I have two responses. First, there are plausible ways for non-substantivists to accommodate the counterfactual robustness of moral principles. For example, neo-Humean accounts of moral principles would explain the regularity of moral principles in virtue of the fact that we inhabit a "lucky" (or "kind") world (as per Lewis, 1994). Namely, we inhabit a world where the distribution of fundamental properties (i.e., the properties at the so-called Humean mosaic) is appropriately organized. Or perhaps such an account can be supplemented with the view that moral properties exhibit a high degree of *elite-ness* (see Dunaway and McPherson (2016)), thus securing their projectability.<sup>11</sup>

Secondly, and more importantly, it seems that the focus on the *metaphysics* of moral principles is fundamentally misplaced, as far as their predictive function is concerned. If it is the case that a moral principle P holds at w, it would not matter whether P is understood substantively or non-substantively: what matters is whether it *holds* at w. And if it *does* hold at w, then we can appeal to P to make predictions (e.g., about Katia's killing). In other words, what matters for prediction is P's *existence* and its *scope*. Having a metaphysical story *explaining* P's nature and existence at w is an interesting but largely orthogonal question.

This is partly why I am not suggesting that **UNIFICATION** can be delivered *only* by non-substantivist views. Naturally, substantivist moral principles can also

<sup>11</sup> The challenge just presented should not be confused with a similar Humean-inspired epistemological objection. As the worry would go, it could be that the actual world *looks as if* it is a world where P holds but, in actuality, regularity breaks down revealing some sort of covert moral particularism. Perhaps Tony's and Katia's killings are morally wrong but Janice's is not as such even though it is similar in circumstances to Tony and Katia's. In response, I agree that (neo-)Humeans have to face this sort of worry (and I cannot fully engage with this argument here) but as a challenge against moral generalizations having a unificatory function it does not hold must sway. *If* a given moral generalization holds in the actual world (never mind how one can know this to be the case), then one can appeal to that principle to make predictions.

unify (and, thus, predict).<sup>12</sup> All I am trying to do in this paper is to say that even if one is a non-substantivist about moral principles, they can still use these principles to do explanatory work.

## 6 Rival Views

Recall that according to my view, moral principles qua moral generalizations are genuinely explanatory given that they explain their instances by subsumption/unification. To compare, other views that take moral principles to be moral generalizations reject this. For example, Berker (2018: 2) suggests that moral principles qua moral generalizations are merely “explanation-involving” rather than “explanation-serving.” In other words, moral generalizations merely *report on* specific explanatory connections between moral and non-moral instances, instead of *themselves* serving as putative explanantia.

Given various assumptions about metaphysical grounding, Berker argues that the only plausible construal of moral principles is in the form of mere moral generalizations. Specifically, this is a view that:

[...] uses no materials other than a wide-scope necessity operator, standard quantification, mundane indicative conditionals, and the full grounding relation, like so:

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

In this sense, Berker effectively bites the bullet and accepts that moral principles qua generalizations, as per the non-substantivist view, “do not themselves partially explain particular moral facts, nor do they explain, govern, or in some other way mediate such first-order explanations.” (Berker, 2018: 26) (i.e., they are not “explanation-serving”). Instead, moral principles are “explanation-involving, so that the features they cite in their embedded conditionals or biconditionals are the grounds of the moral properties at issue.” (Berker, 2018: 26).

Other accounts in the literature have argued explicitly against this view but attempting to identify ways in which moral principles are genuinely explanatory, or, in Berker’s terms, “explanation-serving” (in the sense that they themselves figure in the explanation of particular moral instances). Such accounts (e.g., Enoch (2019); Rosen (2017a, b); Haderlie and Litland forthcoming) adopt a substantivist understanding of moral principles according to which moral principles

<sup>12</sup> Similarly, there could even be particularist ways of predicting by adopting various theoretical epicycles. For example, an anonymous referee suggests that so-called *weak* particularism (via appeal to default or “presumptive” reasons) can deliver the virtue of prediction as stated (see Cullity (2002)). I have my doubts about this view (for critical discussion, see Väyrynen, 2004) but, at any rate, even if successful it would not imply that non-substantive moral principles cannot explain (which is what I am arguing in this paper).

are separate and distinct from their instances and, in turn, *ground* those instances (see also Berker (2018: sec. 6)).

For example, Enoch (2019) argues that in cases where the relevant relata are “just too different” from one another, a (substantivist) moral principle is required in order to bridge the two domains by establishing a connection of metaphysical determination. Additionally, Enoch suggests that there are (at least) two distinct notions of metaphysical grounding at play: metaphysical and normative grounding. To illustrate: moral principles *metaphysically* ground the fact that (say) an action is morally good because it is pleasurable, and the fact that an action is morally good is *normatively* grounded by the fact that it is pleasurable.

Adopting such “grounding pluralism” seems to avoid the charge that moral principles are explanatorily redundant: such principles make it the case that a given act of moral goodness is fully normatively grounded by the fact that it is pleasurable.<sup>13</sup> In other words, even though (internally to the moral context) the goodness of an act is *normatively fully grounded* by its relevant non-moral characteristics, the fact that this is the case is *metaphysically grounded* by a moral principle (for objections, see Berker (2018: sec. 5)).<sup>14</sup>

The approach I am putting forward in this paper marks a distinct way of making sense of the explanatoriness of moral principles. Berker’s approach construes moral principles as moral generalizations but rejects the idea that moral generalizations are genuinely explanatory (i.e., “explanation-serving”). Enoch, on the other hand, accepts that moral principles are explanatory but rejects that they should be understood as moral generalizations. Instead, Enoch (2019) and others (Robinson, 2007, 2011; Rosen, 2017a, b; Haderlie & Litland forthcoming) adopt a substantivist view of moral principles.

In this paper, contrary to Enoch (2019) and Berker (2018), I have argued that moral principles are *both* moral generalizations and genuinely explanatory. Specifically, I have suggested that via the adoption of a specific sort of moral explanatory pluralism, one can plausibly say that moral principles themselves explain via subsumption (while, at the same, time, citing grounding relations between moral and non-moral instances).

<sup>13</sup> Here is a useful illustration: Enoch argues that moral principles *set-up* particular grounding facts (involving moral facts) in the same way certain *legal* facts make it the case (via metaphysical grounding) that a certain law is in effect (e.g., traffic laws). Still, the fact that I did something illegal is fully *legally grounded* in (say) the fact that I broke the speed limit. In other cases, Enoch suggests that non-substantivist laws (what he dubs “Humean laws”) are enough (e.g., concerning principles about conjunctions being grounded in their conjuncts).

<sup>14</sup> Enoch’s grounding pluralism seems importantly different from the explanatory pluralism I adopt in this paper (although it is, no doubt, similar in spirit even though it is used for different purposes). For example, the astute reader should have noticed that I focus on explanations as *vessels* of reporting metaphysical determination relations (rather than taking grounding and explanation to be one and the same), thus effectively adopting the *separatist* view concerning the relation between grounding and metaphysical explanation (Stamatiadis-Bréhier, 2021). Also, I individuate types of explanations in terms of the *project* (and its aims) in which such explanations figure as per **PLURALISM**. Finally, I do not take multiple grounding relations to be primitive in Enoch’s sense (although I do have some sympathies to the view that their primitiveness does not render them necessarily metaphysically mysterious).

## 7 The Circularity Problem

Recall that, according to **PLURALISM**, there are at least two different kinds of moral explanation projects: one falling under **NON-IDLENESS** and one falling under **UNIFICATION**.<sup>15</sup> In this section, I anticipate an objection that potentially arises once we accept these two different types of explanation: what I will call the *circularity problem* (Rosen, 2017a, b; see also Morton, 2020: sec. 4.2.). Roughly, the worry that if one takes moral generalizations to explain their instances, and given that such generalizations are plausibly grounded in those instances, then one is committed to the view that moral generalizations are explanatorily circular. And, as the worry would go, this seems like a *reductio* for either the view that moral principles are generalizations, or **PLURALISM**. In response, I will argue that **PLURALISM** can be used once more to dispel the circularity problem (in the same way it was used to dispel the redundancy challenge).

Consider again the moral generalization that killing is wrong in virtue of its non-moral features. This generalization has certain parts: its instances. In this respect, in a certain context, we can explain this generalization, *constitutively*, in terms of those instances:

**CONSTITUTION:** A generalization P is partially explained by its instance Q.

On the other hand, in another context, we might explain a particular instance of that generalization (for example, the wrongness of an act) in terms of the whole under which it is subsumed:

**SUBSUMPTION:** An instance Q is partially explained by generalization P.

Also, it is plausible that explanations are constrained by the following formal features:

**TRANSITIVITY:** If A explains B, and B explains C, then A explains C.

**IRREFLEXIVITY:** P cannot explain itself.

The circularity problem builds on the fact that these theses cannot all be true at the same time. If the generalizations are explained by their instances, then one cannot appeal to these generalizations to explain those very instances. Doing so would violate either **IRREFLEXIVITY**, or **TRANSITIVITY**, or both. Also, denying either **CONSTITUTION** or **SUBSUMPTION** would go against **PLURALISM**.

Is **TRANSITIVITY** true? Many philosophers think so. But, as noted in the previous section, explanantia must be explanatorily relevant to the explanandum, where explanatory relevance can be cashed out in many different ways depending on the explanatory project under which we operate. Still, there seems to be something true about the idea that certain explanations can *chain* in certain contexts. Is there a way

<sup>15</sup> And, given the nature of **PLURALISM** there could definitely be others as well. A potentially fruitful question for future research would be to further develop **PLURALISM** and identify even more functions that can be served by distinct explanatory projects.



to reconcile something in the vicinity of **TRANSITIVITY** with the pluralist picture I have been proposing?

The natural way to proceed is to make **TRANSITIVITY** more fine-grained. As already noted, there are many ways to cash out the notion of explanatory relevance. **SUBSUMPTION** can be understood as employing a constraint like **UNIFICATION**. This constraint makes sure that the information that is eventually included in the explanation contributes towards the epistemic goal of unifying the relevant phenomena. It is obvious that this constraint does not exclude transitive explanations. For example, a more fleshed out moral version of **SUBSUMPTION** might be the following:

**SUBSUMPTION\***: The fact that a particular killing is morally wrong is partially explained by the fact that the principle “you ought not to kill” holds. Now assume, for simplicity, that every killing minimizes utility. The fact that the principle “you ought not to kill” holds is, in turn, partially explained by the fact that the principle of utility is true. Thus, the fact that this particular killing is morally wrong is, transitively, explained by the fact that the principle of utility holds.

I take **SUBSUMPTION\*** to be uncontroversial insofar as we can make sense of an explanation that proceeds by unifying. In this sense, I assume that derivative principles can be explained and do, themselves, explanatory work. The wrongness of a particular killing is first incorporated under a larger regularity (“you ought not to kill”), which, in turn, is subsumed under an even more general regularity (the principle of utility).

Both explanations figuring in **SUBSUMPTION\*** have the same *aim*. **SUBSUMPTION\***, in this sense, provides a hint towards the right view of transitivity. **TRANSITIVITY\***: If A explains<sub>1</sub> B and B explains<sub>2</sub> C, then A explains C, iff, both explanations accommodate the aims and goals of the research project in which they figure.

Now consider the two following explanations:

**M**: [Killing is morally wrong] partially explains (via subsumption) [This particular killing is morally wrong].

**M+**: [This particular killing is morally wrong] partially explains (constitutively) [Killing is morally wrong], and, [Killing is morally wrong] partially explains (via subsumption) [This particular killing is morally wrong]. Thus, [This particular killing is morally wrong] partially explains itself.

**M** is unobjectionable whereas **M+** is the circular result that the circularity problem warns us about.

At this point, we should immediately notice that **M+** is **M** plus an extra component. **M+** involves the explanation of a particular killing by subsuming it under

a moral regularity, but then it proceeds by explaining that very principle constitutively. But this additional step should be resisted. Knowing the constituents of a principle does not help the unificationist goal of **M**. In this sense, **M+** is inferior to **M** since it introduces irrelevant information. Thus, by **TRANSITIVITY\***, we can avoid the circularity worry without giving up on a sufficiently substantive version of the transitivity principle.<sup>16</sup>

One assumption in the previous argument is that **M+** has the same aim as **M**. But this is not necessarily so. One might claim that there are contexts where explanations like **M+** are *exactly* what we should be looking for. It is obvious that we can manufacture an explanatory project with aims **S**, where it is stipulated that **S** requires **SUBSUMPTION** and **CONSTITUTION** to be linked. But to my mind, these cases are atypical for they require an implausible and highly artificial explanatory project. Why would someone be interested in an explanation like **M+**? In this sense, I am happy to grant that there are atypical cases where explanations are circular.

Nor do I wish to suggest a general ban on explanations involving both unificationist and constitutive goals. For example, consider the following explanation:

**MIXED:** [Killing is morally wrong] partially explains (via subsumption) [This particular killing is morally wrong] and, in turn, [This particular killing is morally wrong] partially explains (constitutively) [This particular killing is morally wrong or imprudent]. Thus, [Killing is morally wrong] partially explains [This particular killing is morally wrong or imprudent]

But accepting the potential legitimacy of **MIXED** is compatible with the rejection of circular explanations like **M+**. In other words, I am not appealing to **PLURALISM** to argue that mixed explanations are *in general* problematic. For all we know, multiple explanatory relevance constraints are systematically mixed and matched. What I am saying, rather, is that *circular* mixed explanations like **M+** are plausibly ruled out with the help of **PLURALISM**. It is just hard to see what legitimate research project would require the combination of unificationist and constitutive relevance norms *in this particular way*. Until the objector provides a plausible case where this is so, the circularity challenge does not pose a threat to the explanatoriness of moral generalizations.

## 8 Conclusion

In this paper, I defended the explanatoriness of moral principles qua moral generalizations. I argued that even if moral generalizations are explanatorily redundant under **NON-IDLESNESS**, they are not so under **UNIFICATION**. In other words, one can appeal to moral principles to explain a particular instance falling within its scope via *subsumption*. Then, I defended **UNIFICATION** by situating it in the

<sup>16</sup> Analogous strategies have been appealed to in the literature on neo-Humean scientific laws (for discussion, see Bhogal (2020: sec. 2.4.)) (although this is the first time, as far as I can tell, that a principle like **TRANSITIVITY** is applied to the moral domain).

context of what I take to be the independently plausibly theory of moral explanatory pluralism. Specifically, **UNIFICATION** can be used in explanatory projects which aim at further *understanding* of the good, as well as in terms of *prediction*.

Finally, I tackled the circularity objection. Moral generalizations explaining their instances via subsumption seems to be in tension with the idea that moral generalizations are grounded in terms of those very instances. Under apparently plausible assumptions about the transitivity and irreflexivity of explanation, it could be objected that, on my picture, moral principles end up explaining themselves. As mentioned, one could argue that if under the pluralist picture moral principles are explanatorily circular, then this is a kind of *reductio* either for the pluralist view or the view that moral generalizations are genuinely explanatory.

In response, I argued that by appeal to the pluralist machinery once more, one can put forward a plausible *amendment* of the transitivity principle thus effectively blocking the result that moral generalizations explain themselves (except, perhaps, in highly arbitrary and philosophically uninteresting cases).

This further suggests that, for all intended purposes, moral generalizations can in fact explain. Of course, this does not settle the debate between substantivism and non-substantivism about moral principles, but it makes some non-trivial headway.

**Acknowledgements** For written comments on previous and related drafts I thank: Pekka Väyrynen, Jack Woods, Gideon Rosen, Robbie Williams, David Kovacs, and a very helpful anonymous reviewer from *Acta Analytica* (as well as reviewers from other journals). I also thank audiences from Sheffield and Leeds. A somewhat distant precursor of this paper (with a significantly different dialectic) was titled 'How moral generalizations explain'.

**Funding** Open access funding provided by Tel Aviv University. This work was supported by the Azrieli Foundation via an International Postdoctoral Fellowship.

## Declarations

**Conflict of Interest** The author declares no competing interests.

**Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

## References

- Baker, D. (2021). If you're quasi-explaining, you're quasi-losing. In Russ Shafer-Landau (ed.), *Oxford Studies in Metaethics 16*. Oxford University Press.
- Berker, S. (2018). The explanatory ambitions of moral principles. *Noûs*, 53(4), 904–936.
- Bhogal, H. (2017). Minimal anti-Humeanism. *Australasian Journal of Philosophy*, 95(3), 447–460.

- Bhagal, H. (2020). Humeanism about laws of nature. *Philosophy Compass*, 15(8), 1–10.
- Boyd, R. (1988). How to be a moral realist. In G. Sayre-McCord (Ed.), *Essays on moral realism* (pp. 181–228). Cornell University Press.
- Boyd, R. (1999). Kinds, complexity, and multiple realization. *Philosophical Studies*, 95(1–2), 67–98.
- Brink, D. O. (2013). *Mill's progressive principles*. Oxford University Press UK.
- Cartwright, N., Pemberton, J., & Wieten, S. (2020). Mechanisms, laws and explanation. *European Journal for Philosophy of Science*, 10(3), 1–19.
- Craver, C. F. (2009). *Explaining the brain*. Oxford University Press.
- Cullity, G. (2002). Particularism and moral theory: Particularism and presumptive reasons: Garrett Cullity. *Aristotelian Society Supplementary*, 76(1), 169–190.
- Dorsey, J. E. (2016). On the grounding-reduction link. *American Philosophical Quarterly*, 53(4), 411–422.
- Fogal, D., & Risberg, O. (2020). The metaphysics of moral explanations. *Oxford studies in metaethics* (vol. 15, pp. 170–194).
- Dunaway, B., & McPherson, T. (2016). Reference magnetism as a solution to the Moral Twin Earth problem. *Ergo: An Open Access Journal of Philosophy* 3.
- Enoch, D. (2019). How principles ground. In *Oxford Studies in Metaethics* 14:1–22. Oxford: Oxford University Press.
- Giannotti, J. (2022). Fundamentality and minimalist grounding laws. *Philosophical Studies*, 179(9), 2993–3017.
- Haderlie, D., C. & Litland, J., E. (forthcoming). Grounding Legalism. *Philosophical Quarterly*.
- Kitcher, P. (1989). Explanatory unification and the causal structure of the world. In P. Kitcher & W. Salmon (Eds.), *Scientific Explanation* (pp. 410–505). University of Minnesota Press.
- Kovacs, D. M. (2020a). The oldest solution to the circularity problem for Humeanism about the laws of nature. *Synthese*, 198(9), 1–21.
- Kovacs, D. M. (2020b). Metaphysically explanatory unification. *Philosophical Studies*, 177(6), 1659–1683.
- Kovacs, D. M. (2021). There is no distinctively semantic circularity objection to Humean laws. *Canadian Journal of Philosophy*, 51(4), 270–281.
- Krämer, S., & Roski, S. (2017). Difference-making grounds. *Philosophical Studies*, 174(5), 1191–1215.
- Lewis, D. (1983). New work for a theory of universals. *Australasian Journal of Philosophy*, 61(4), 343–377.
- Lewis, D. (1994). Humean supervenience debugged. *Mind*, 103(412), 473–490.
- Litland, J. E. (2013). On some counterexamples to the transitivity of grounding. *Essays in Philosophy*, 14(1), 3.
- Machamer, P., Darden, L., & Craver, C. F. (2000). Thinking about mechanisms. *Philosophy of Science*, 67(1), 1–25.
- Makin, M. (2019). Rigid/non-rigid grounding and transitivity. *Inquiry: An Interdisciplinary Journal of Philosophy*, 62(2), 136–150.
- Marshall, D. (2015). Humean laws and explanation. *Philosophical Studies*, 172(12), 3145–3165.
- Maudlin, T. (2007). *The metaphysics within physics*. Oxford University Press.
- McPherson, T. (2012). Ethical non-Naturalism and the metaphysics of supervenience. *Oxford Studies in Metaethics* 7.
- Morton, J. (2020). Grounding the normative: A problem for structured non-naturalism. *Philosophical Studies*, 177(1), 173–196.
- Potochnik, A. (2015). Causal patterns and adequate explanations. *Philosophical Studies*, 172(5), 1163–1182.
- Railton, P. (1980). *Explaining explanation: A realist account of scientific explanation and understanding*. Dissertation, Princeton University
- Ridge, M. (2007). Anti-reductionism and supervenience. *Journal of Moral Philosophy*, 4(3), 330–348.
- Robinson, L. (2007). Moral principles are not moral laws. *Journal of Ethics and Social Philosophy*, 2(3), 1–22.
- Robinson, L. (2011). Moral principles as moral dispositions. *Philosophical Studies*, 156(2), 289–309.
- Rodriguez-Pereyra, G. (2006). Resemblance nominalism: A solution to the problem of universals. *Philosophy and Phenomenological Research*, 72(1), 241–246.
- Rosen, G. (2017). Metaphysical relations in metaethics. In T. Routledge (Ed.), *Tristram Colin McPherson & David Plunkett* (pp. 151–169). Routledge.
- Rosen, G. (2017a). What is a moral law? *Oxford Studies in Metaethics* 12.
- Ross, D. (1939). *Foundations of ethics*. Oxford University Press.
- Sider, T. (2020). Ground grounded. *Philosophical Studies*, 177(3), 747–767.
- Sinclair, N. (2024). Metaethics and the nature of properties. *Aristotelian Society Supplementary*, 98(1), 133–152.

- Stamatiadis-Bréhier, A. (2021). Backing as truthmaking. *Canadian Journal of Philosophy*, 51(5), 367–383.
- Stamatiadis-Bréhier, A. (2022). Nomic moral naturalness. *Inquiry*, 1–22.
- Stamatiadis-Bréhier, A. (2023). Grounding functionalism and explanatory unificationism. *Journal of the American Philosophical Association*, 9(4), 799–819.
- Strevens, M. (2008). *Depth: An account of scientific explanation*. Harvard University Press.
- Suikkanen, J. (2024). Metaethics and the nature of properties. *Aristotelian Society Supplementary*, 98(1), 113–131.
- Suikkanen, J. (2024). Nonnaturalism, the supervenience challenge, higher-order properties, and trope theory. *Journal of Ethics and Social Philosophy*, 26(3), 601–632.
- Väyrynen, P. (2004). Particularism and default reasons. *Ethical Theory and Moral Practice*, 7(1), 53–79.
- Väyrynen, P. (2023). Moral generalism and moral particularism (2nd edition). In Christian B. Miller (ed.), *The Bloomsbury Handbook of Ethics*. Bloomsbury Academic. pp. 381–396.
- Walden, K. (2016). The relativity of ethical explanation. *Oxford Studies in Normative Ethics*, 6.
- Woods, J. (2018). Emptying a paradox of ground. *Journal of Philosophical Logic*, 47(4), 631–648.
- Woodward, J., & Hitchcock, C. (2003). Explanatory generalizations, part I: A counterfactual account. *Noûs*, 37(1), 1–24.
- Woodward, J. (2003). *Making things happen: A theory of causal explanation*. Oxford University Press.
- Wygoda Cohen, S. (2020). Not all partial grounds partly ground: Some useful distinctions in the theory of grounding. *Philosophy and Phenomenological Research*, 100(1), 75–92.

**Publisher's Note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.