The Projectability Challenge to Moral Naturalism

[Terence Cuneo & Andrew Reisner: Draft of 13 June 2019]

1. Introduction

Suppose you know that killing, lying, stealing, torturing, maiming, and the like are morally wrong. You also know that you're required to treat others with fairness and honesty and your family with love. Still, there are many moral issues to which you've given almost no attention, such as the situation you encounter in:

> **Forgiveness**: An old enemy of your family, whose actions have caused your family considerable hardship and pain, has approached you, asking for forgiveness. By all appearances, he is sincere. You have, however, never encountered a situation quite like this and your ethical worldview has little to say about which moral concepts apply to repentant enemies. Should you ignore his request? Continue to hold him accountable? Require reparation? Or are you morally required to forgive him? After reflecting on the case, you determine that while there is no general requirement to forgive, it is nevertheless right to do so in this case, as it is for the best for both you and him. In this case, forgiving your repentant enemy falls under the moral concept RIGHT.

Other situations we encounter are different, requiring us to revise our ethical understanding in light of our experience. One such situation is the focus of:

> **Spying**: Recently, your cousin has revealed to you that he has been spying on his neighbor, invading his privacy. He is committed to doing nothing with any information

1

he may gain, so no one is harmed. For years, however, you have held that the moral

concept WRONG applies to only what causes harm. But upon hearing your cousin's

confession, you come to hold that it applies to "victimless" actions too, such as when

someone spies on another for purely prurient reasons, invading his privacy. You

experience a shift in your ethical convictions, expanding your understanding of a central

ethical concept WRONG, which you subsequently apply in this case.[1]

Cases such as Forgiveness and Spying help us to appreciate that the ethical life has twists and

turns, often calling for us to make ethical determinations that are unfamiliar in various respects.

We call moral situations of this sort *novel*.

Novel moral situations (or simply 'novel situations') come in at least two sorts. Those of

the first sort, such as Forgiveness, involve a *novel case*. That is, (i) they involve a case of a type

that you have not encountered previously in which an action has morally relevant features, but

(ii) your ethical worldview does not situate you to immediately ascertain what contribution these

features make to that action's moral status. Yet (iii) on the basis of experience or reflection, you

size up the case, applying a moral concept to it.[2] In contrast, a situation of the second sort, such

as Spying, involves a *novel application* of a moral concept. A situation of this sort (i) may

involve a type of case with which you are familiar but (ii) its moral dimensions are such that, due

to your ethical worldview, you either previously failed to recognize them or took them to be

rather different from how they now strike you. Still, (iii) upon experiencing or reflecting on its

---

[1] We rely on the assumption that not all wrongs are harms. If you reject this assumption, we invite you to interpret the case in such a way that you begin by holding that, when a victim suffers harm of the relevant sort, the concept WRONG applies to actions that cause such harm only when the harm is known. But reflection on your cousin's confession enables you to understand that, when a victim suffers harm of the relevant sort, the concept WRONG also applies to actions that cause such harm even when the harm is unknown (by anyone else other than the wrongdoer).
[2] Throughout we use 'experience' broadly, to cover mental states or events that are (*inter alia*) perceptual, affective, imaginative, or cognitive.

morally relevant features, you revise your understanding of a relevant moral concept, applying it to the case at hand.

Let's say that if you were to make an accurate and justified moral judgment in a novel situation of either sort by applying a moral concept, then you have *projected* that concept (or 'projected,' for short).[3] Novel moral situations introduce a datum that we'll call:

Projectability: We can, on the basis of experience or reflection, accurately and justifiedly apply moral concepts in novel situations.

The task of explaining Projectability is what we call the *projectability challenge.*

In our view, every metaethical theory must address this challenge.[4] Here we focus on realist versions of moral naturalism and nonnonaturalism, understanding these views to concern the nature of moral properties and facts. While both types of view hold that there are moral properties and facts, moral naturalists assert that it belongs to their nature to be natural; nonnaturalists deny this. Our project is to explore how views of these sorts can answer the projectability challenge, arguing for two primary claims. First, paradigmatic nonnaturalism can satisfactorily answer the challenge, explaining how agents can project. This explanation, as we'll emphasize, has both the right form and incorporates the correct normative content. Second, it is unclear whether naturalist views can satisfactorily answer the challenge. The conclusion we draw is not that moral nonnaturalism is true or likely to be true. Rather, it is that there is an important respect in which paradigmatic nonnaturalism holds an advantage over its most prominent

---

[3] Our use of the term "projection" is related to but distinct from discussions of projection in epistemology and philosophy of science, where what is at issue is, for example, the use of predicates in induction and lawlike generalizations (e.g., whether to use "blue" and "green" or "grue" and "bleen").

[4] Or at least adequately explain why they need not do so. In §6, we revisit the issue, acknowledging that analogues to the projectability challenge apply to a wide range of domains beyond metaethics.

naturalist rivals. The conclusion is interesting if only because it is widely assumed that naturalism has an easier time handling thorny problems in moral epistemology. We argue that there is at least one such problem of which this assumption is not true.

We frame our discussion around the following question: What understanding of moral thought and reality is well-suited to explain how we can project? We maintain that it is one according to which moral concepts and properties have what we call an *essence-profile*. Such a profile is (roughly) a range of substantive truths that state not only what a concept or property is at its very core, but also includes truths that are necessarily grounded in what lies at its core (a subset of these truths are what the medievals called *propria*). This understanding of concepts and properties, we contend, enables us to see how moral concepts could bear the appropriate sorts of relations to their corresponding moral properties such that we could project. We further maintain that this understanding not only fits comfortably with a non-naturalist metaethic, but also draws upon an important strand within this tradition. The same is not true, we claim, of moral naturalism, as these views either fail to provide explanations of the right form or there is insufficient reason to believe that they can incorporate the correct normative content. But we stop short of claiming that naturalists cannot answer the projectability challenge, claiming only that doing so will probably require them to take their views in new directions, enriching them in ways that may not easily mesh with their fundamental commitments.

2. Clarifying the challenge
To bring the projectability challenge into sharper focus, let us offer five preliminary remarks regarding how we understand it and what a satisfactory answer requires.

4

First, answering the projectability challenge requires explaining not how ideally rational agents could project, but how, in principle, agents like us with our capacities in our circumstances could do so. For our purposes, the differences between ideal and ordinary agents is significant, since idealized agents are typically specified as having complete information, unlimited cognitive resources, and perfect inferential abilities. Ordinary agents, however, must make do with much less. If they can project, it will have to be on the basis of limited information by way of the employment of limited cognitive, informational, and inferential capacities. In circumscribing the task in this way, we make no assumption that we're excellent at projecting moral concepts. Still, we do seem to be able to do it, and we seem good enough in a wide enough range of cases for projection to be something that needs explaining.[5]

Second, answering the projectability challenge requires accounting for the scope and epistemic status of projection. What is at issue is not simply how we apply moral concepts to novel situations but also how we can *accurately* and *justifiedly* apply these concepts in novel situations, where such application could underwrite knowledge attributions to those who project. We thus propose to set aside cases in which agents project in a way that is accurate simply because of luck or mere happenstance.

This restricts what could count as an explanation of projection. Suppose, for example, one were to appeal to some familiar heuristic used in moral decision making, such as a version of

---

[5] Jackson, Pettit, and Smith (2000) represents the only other discussion of projection in the moral domain of which we're aware. They contend that, while non-naturalism is ill-equipped to handle (something similar to) the projectability challenge, naturalism is. Here we bracket discussion of their paper because its solution to (something like) the projectability challenge assumes that there are patterns among natural features that *ideal* agents would discern (90-2). Even if there were such patterns, this would not explain how agents like us, who are not ideal, can project. It might be that Jackson, Pettit, and Smith are assuming that moral experience helps to bring our actual dispositions to apply moral concepts in certain ways close to those of idealized moral agents. But absent an explanation of how that could be so (and why an appeal to ideality contributes to an explanation), we doubt that this claim can serve as the basis for an answer to the projectability challenge. It cannot simply be assumed that our actual responses track those of an ideal agent.

the Golden Rule (roughly, "Do unto others as you would have them do unto you"). While the use of such heuristics has value, especially in moral education, its value is limited inasmuch as it lacks the requisite alethic and epistemic credentials: it is not sufficiently trustworthy, as the fact that having formed a judgment on the basis of the Golden Rule does not by itself warrant, let alone entail, knowledge attributions. These points apply mutatis mutandis to other heuristics that could be reasonably ascribed to ordinary moral agents (e.g., "Act as Jesus would").

Third, answering the projectability challenge requires accounting for the role of experience and reflection in projection. When we accurately and justifiedly apply moral concepts to novel situations, we do not do so willy-nilly but on the basis of how things strike us morally—such as when it strikes you that actions are sometimes wrong even when they harm no one.

How things strike us morally is often a matter (at least to some extent) of discerning patterns among the natural features upon which moral features supervene. In Spying, for example, you change your mind about the nature of rightness by becoming vividly aware of the ways in which your cousin has treated his neighbor: he fails to respect his neighbor's right to privacy but does not harm him. If this is correct, any attempt to explain Projectability will have to offer a view about whether there are patterns that we apprehend at the subvenient natural level and how we can project moral concepts on the basis of our apprehension of these patterns.

This brings us to our fourth preliminary remark. We assume that a viable view of the indicated patterns should respect the "shapelessness" of the natural with respect to the moral. What we mean by this is that, in a wide range of cases, the natural properties (or a subset of such properties) that serve as the supervenience bases for moral properties are too numerous,

6

complex, heterogeneous, and variable in their normative valences to be accurately categorized

with natural concepts alone (at least from the deliberative perspective of non-ideal agents).[6]

To illustrate, consider the following range of act-types, specified in purely natural terms,

which are central cases of wrongdoing:

killing others for sport;

cheating on one's tax forms;

being indifferent to others' flourishing;

dissembling for reputation's sake;

snubbing others because of their physical appearance;

ridiculing others due to their flaws;

abusing one's position of power; and so on.

Reflecting on these and other actions that could be wrong prompts the question: Is there some

pattern that is both describable using natural concepts alone and discernable from the

deliberative perspective of a non-ideal agent that would be present in each and every case in

which—and only in those cases in which—these actions are wrong?

Not obviously. Although the issues in question are vexed, the act-types cited above

represent only a tiny sample of the range of actions that can be wrong. Were we to expand the

list, we suspect that the prospects of discovering a pattern of the right sort would probably be

much more difficult than the examples offered above would suggest. What is more, as moral

holists have emphasized, subvenient natural features often switch their normative valences. In

one case, bowing to a superior might be a sign of respect; in another, it might be a deep insult. In

---

[6] We take this last claim to be compatible with the view that moral properties are themselves natural. In that case, the claim would be that, in a wide range of cases, there is some subset S of natural properties on which a moral property M supervenes that is not identical with M, but the members of S are too numerous, complex, heterogeneous, and variable in their normative valences to be accurately categorized with natural concepts alone (at least from the deliberative perspective of non-ideal agents).

one case, telling falsehoods might be wrong; in another, it might be required. And so on.[7] Taking

full account of the ways in which the moral status of actions can vary with context would be

unmanageable, at least for beings with our sorts of cognitive and practical limitations. It is

because the natural properties on which moral properties supervene are both vast (perhaps

innumerable) and very different (perhaps wildly heterogeneous) that, when we ask what they

have in common, often the only answer is to group them under moral concepts, such as

WRONGING ANOTHER.[8]

The thesis that the natural is shapeless in the way just described is important because it

gives content to the idea that answering the projectability challenge is not trivial. For if

subvenient natural features are shapeless with regard to the moral, then addressing this challenge

cannot be a matter of simply appealing to patterns of such features discernable from the

deliberative perspective by a non-ideal agent using only natural concepts. Or, to put the point

more modestly: while there may be such patterns, in a wide range of cases, there appear to be

none that are sufficiently general and informative that can form the basis from which agents such

as us with our capacities in our circumstances using only natural concepts can project.

Fifth, we'll assume that, while a satisfactory answer to the projectability challenge must

explain how agents can project, we envision such an explanation as being limited in two

important respects. For one thing, it needn't be a full explanation of how agents can project, one

that (say) enlists a suite of psychological processes or mechanisms whose workings enable

agents to project.[9] So long as we have the fundamental building blocks of an explanation in

place, which can provide an explanation of the right form and with the correct normative content,

---

[7] See Dancy (2004, Part III).
[8] Oddie (2005, 144) makes a similar point.
[9] We imagine that there are a variety of promising proposals regarding what the psychological processes or mechanisms are. For one possible example, *cf.* Railton (2017).

we can allow that a deeper and more complete explanation would add further detail and nuance. Indeed, we're aware that the explanation we provide below goes no further than identifying these fundamental building blocks. Second, we'll assume that a satisfactory answer to the projectability challenge must account for how agents (of the requisite sort) are *in a position* to project.[10] While we have no exact characterization of what it is for such an agent to be in a position to project, the idea is that an agent is in such a position if she would project provided that enough goes well with her. For example, elements of her environment do not interfere with grasping the situation before her, her faculties work well enough, she notices relevant features of her situation that we would expect her to, and so forth.

These five preliminary remarks help us to characterize what a satisfactory answer to the projectability challenge would be, at least for positions with moral realist commitments.

3. Paradigm nonnaturalism

In his discussion of Open Question-style arguments, Philip Stratton-Lake writes that nonnaturalists have assumed that "there is a certain isomorphism between the structure of our concepts and the nature of the world, such that a proper analysis of our concepts would reveal to us the nature of the corresponding property or thing."[11] We begin this section by sketching the *traditional view of concepts and properties* (or, simply the 'traditional view') that makes sense of this assumption, so widely embraced by nonnaturalists.[12] We then present a view that combines the traditional view with moral nonnaturalism. We call it *paradigm nonnaturalism*, arguing that it can answer the projectability challenge.

---

[10] §4 touches on what it is to have reasonable mastery of a concept.
[11] Stratton-Lake (2002, 9).
[12] [citation omitted] develops the position in more detail, including the rationale for what we call (below) Isomorphism.

Three claims of increasing substance constitute what we're calling the traditional view. The first, which concerns the relationship between concepts and properties, states that concepts are referential devices or means of getting properties in mind. Any property to which a concept refers is its *corresponding property*.

The second claim, which concerns our grasp of concepts, says that concepts are not only graspable in the sense of being the possible objects of mental relations such as entertaining, understanding, and learning, but also that we can grasp them in different ways and to different degrees.[13] One's grasp of a concept may be better or worse, along (at least) two dimensions. First, it may be incomplete, as when the irrationality of π is not yet within one's purview. Second, it may be incorrect, as when one is such as to take π to be rational. At one extreme lies perfect mastery of a concept; at the opposite extreme, just short of complete lack of a concept, lies mere possession of it. In between lies *reasonable mastery*, which is the kind of competence with a concept that one possesses when one grasps its basic or paradigm applications. We assume that agents often have this sort of competence with many moral concepts.[14]

The third and most substantial claim is that both concepts and properties have essence-profiles, which specify what it is to be a given concept or property. A concept's essence-profile includes (but is not exhausted by) its application conditions. We think of these as being given by a concept's *application conditionals*, which identify some entity to which a concept might apply, specifying whether that concept applies to that entity.[15] For example, the traditional view would say the following regarding the concept WRONG.

---

[13] It is a substantive question what it is to grasp a concept and, hence, what it is to grasp (what we below call) the application conditionals that compose its essence-profile. *Cp*. Peacocke (1992), Bealer (1998, Part II), Chalmers and Jackson (2001), and Peacocke (2009). Our discussion is compatible with each of these (and other possible) answers.

[14] [Citation omitted] provides a fuller treatment of what reasonable mastery of a concept consists in.

[15] Chalmers and Jackson (2001) and Korman (2006) employ this terminology. *Cp*. Peacocke (2008) on "reference rules."

It belongs to the essence-profile of the concept WRONG that:

necessarily, if some action falls under the concept KILLING OTHERS FOR SPORT, then WRONG applies to that action;

necessarily, if some action falls under the concept DISSEMBLING FOR REPUTATION'S SAKE, then WRONG applies to that action;

necessarily, if some action falls under the concept ABUSING ONE'S POSITION OF POWER, then WRONG applies to that action; and so forth.

These application conditionals are *central* in the sense that they specify a set of paradigm applications of the concept WRONG. Anyone with reasonable mastery of the concept would grasp them.

Like concepts, properties also have essence-profiles. A property's essence-profile will include (but not be exhausted by) the conditions under which it is exemplified. As with concepts, we can think of these conditions as given by an array of conditionals, which we'll call a property's *exemplification conditionals*. These identify some entity that might exemplify that property, specifying whether that entity exemplifies that property. For example, the traditional view would say the following about the property *being wrong*.

It belongs to the essence-profile of the property *being wrong* that:

11

necessarily, if some action instantiates *killing others for sport*, then it instantiates *being wrong*.

necessarily, if some action instantiates *dissembling for reputation's sake*, then it instantiates *being wrong*;

necessarily, if some action instantiates *abusing one's position of power*, then it instantiates *being wrong*; and so forth.

These conditionals are *central* in the sense that they specify a set of basic or paradigm exemplifications of wrongness that anyone with reasonable mastery of WRONG would grasp via the grasp of their corresponding application conditionals.[16]

We understand paradigm nonnaturalism to be the view that combines moral nonnaturalism with the traditional view of concepts and properties. That is, we understand the view to conjoin the three claims that constitute the traditional view with the thesis that there are non-natural moral properties and facts. In the remainder of this section, we explain how this view can answer the projectibility challenge.

The answer begins with a claim that should be attractive to a wide range of realist theories, namely:

---

[16] There are other important questions to address about these conditionals. How do concepts enable us to get them in mind? Do they belong to the very essence of WRONG and *wrong*, rather than merely belong to their essence-profiles? Do they track grounding relations, such that what is specified in their antecedents grounds what is specified in their consequents? Are they defeasible? Are they best understood to be the materials of a two-dimensional semantics? While we have views on these matters, we'll not address them here.

Grasp: In a range of cases, agents improve their reasonable mastery of moral concepts in virtue of having experiences as of their corresponding moral properties (or their instantiations).

As noted earlier (see n2), such experiences need not be perceptual; they could be reflective, intuitive, judgmental, affective, imaginative, or of some other kind.

Paradigm nonnaturalism then appeals to a modest extension of the traditional view, connecting an agent's reasonable mastery of moral concepts to the activity of projection:

Insight: In a range of cases, improved mastery of a concept puts one in a position to judge accurately and justifiably in a novel situation whether that property is instantiated.

The leading thought here is that, via improved mastery of a concept, agents are able to move beyond the central cases (or conservative extensions of these cases), which feature in that concept's central application conditionals, applying the concept to novel cases accurately and justifiedly.

The traditional view helps explain how agents can in this way improve their reasonable mastery of moral concepts, as it commits itself to:

Isomorphism: The application conditions of moral concepts are isomorphic with respect to the exemplification conditions of their corresponding moral properties.

To say that the application conditions of a moral concept such as WRONG are isomorphic with respect to the exemplification conditions of its corresponding property is simply to claim that an action specified in the antecedent of its application conditionals falls under the concept WRONG when and only when that action exemplifies the property *being wrong*. This has the following notable epistemic implication: the application conditions of this concept make evident the conditions under which something has property *being wrong*, revealing the conditions under which something has that property.

These three theses, we maintain, enable paradigm nonnaturalism to explain:

Projectability: We can, on the basis of experience or reflection, accurately and justifiably apply moral concepts in novel situations.

To illustrate, return to the case we called Spying. When informed of your cousin's behavior, you undergo a series of reflections involving moral experience as of wrongness (per Grasp). On the basis of this experience, you come to grasp the concept WRONG to a degree and in a way that you did not previously, improving your reasonable mastery of it (per Insight). For, on the basis of this experience, you now realize that:

if something is a case of spying for merely prurient reasons, then it falls under the concept WRONG.

14

This improved understanding of WRONG, in turn, puts you in a position to grasp a subset of the exemplification conditionals of its corresponding moral property *being wrong* (per Isomorphism). For, provided that all goes well, you will see that:

if something is a case of spying for merely prurient reasons, then it is wrong.

In other words, on the basis of your experience, you are in a position to appreciate that actions have the property *being wrong* even when they are "victimless." Since, by hypothesis, you judge that your cousin's behavior is an instance of the relevant action—namely, spying on others for merely prurient reasons—you accurately and justifiedly apply the concept WRONG in this novel situation. For when your judgment flows from reasonable mastery of the concept, which gives you insight into the essence-profile of its corresponding property, your judgment is not a matter of luck but is such as to underwrite an attribution of genuine knowledge (at least when your judgment is accurate).

In this way, on the basis of experience, you gain insight into important dimensions of the concept WRONG and its corresponding property—dimensions which you had to this point failed to grasp. Your previous grasp of this concept and property was at best incomplete, perhaps even distorted in certain respects. Your experience in Spying is transformative in the sense that you gain moral understanding that you had previously lacked, which puts you in a position to project.

Let us call attention to two features of this explanation.

First, this explanation assumes that having reasonable mastery of the moral concept WRONG involves grasping a substantive subset of its application conditions, including those that we designated as central. It follows that, if one has reasonable mastery of this concept, then one

15

is in a position to grasp its corresponding property's exemplification conditions, given the isomorphism between the two sets of conditions. And if one's reasonably mastery of WRONG is enhanced by having an experience as of wrongness in a novel situation, then one is in a position to project, just as Insight says. That one is in such a position is (in part) an implication of paradigm nonnaturalism, for the view implies that these exemplification conditions are not "hidden" behind or within some essence distinct from the essence-profile of its corresponding concept. In this way, paradigm nonnaturalism incorporates the non-naturalist commitment to which Stratton-Lake calls attention, namely, that there is a "certain isomorphism between the structure of our concepts and the nature of the world," such that a proper understanding of the essence-profile of moral concepts would reveal central components of the essence-profile of their corresponding properties.

Second, this explanation assumes that the central application conditionals of WRONG and its corresponding exemplification conditionals have a certain *normative-profile*. The concept WRONG necessarily applies to such actions as killing others for sport, breaking promises for mere convenience's sake, and not abusing a position of power. This is not a superficial feature of paradigm nonnaturalism. For the non-naturalist tradition is plausibly interpreted as being particularly concerned to yield metaethical views that yield the right sort of first-order moral verdicts in central cases. Think, in this regard, of Open Question-style arguments. Under a plausible reading, these arguments draw attention to the fact that certain attempts to identify moral concepts or properties with natural ones feel "open" not (only) because what they propose is incomplete, overly vague, or uninformative. It is rather because what they propose is normatively wrong-headed, yielding characterizations of moral reality that fail to have the right sort of normative-profile. It is on the basis of these first-order judgments that Moore and his

16

followers have concluded that properties such as *failing to enhance evolutionary fitness* could not be what moral wrongness is.[17] This point—which we'll return to later—is significant because paradigm non-naturalism's explanation of Projection turns on the claim that agents with reasonable mastery of the concept WRONG can project because they grasp a concept and property that of necessity have *this* sort of normative-profile.

The paradigm non-naturalist explanatory strategy we've presented maintains that agents sometimes have experience as of moral properties, that they improve their reasonable mastery of moral concepts as a result, and that such improvement positions them to gain insight into important dimensions of those properties, such as those they encounter in novel cases. We've claimed that this explanation has the correct form (in virtue of appealing to Grasp, Insight, and Isomorphism) and incorporates the right content (in virtue of its commitment to moral concepts and features having their particular essence-profiles).

These are not trivial commitments. That we have experiences as of moral reality and have insight into its essence-profile are particularly substantive commitments. Some would charge that paradigm non-naturalism's claim that we have experience as of moral properties sits poorly with its claim that these properties are not part of the natural world, which is characterized in causal terms. Others might worry that the commitment to our being able to have insight into the essence-profile of moral concepts and properties presupposes a mysterious intellectual faculty, an "essence detector," which should be rejected. (Still others may doubt whether the natural is shapeless with respect to the moral.) If you harbor such reservations, then we invite you to interpret the account we've offered conditionally: on the supposition that paradigm nonnaturalism can make sense of the idea that we can have experiences as of moral properties

---

[17] *Cp.* Moore's (1903, §34) comments regarding evolutionary ethics, as well as Hurka (2014, 95). [citation omitted] develops this interpretation of the Open Question Argument.

and that we have the relevant mental capacities, then this is the explanation of projection that it can offer.[18] For what it is worth, we are happy to spot naturalists similarly controversial assumptions. For example, we'll assume below that naturalists can make sense of the idea that infinite or massively disjunctive properties are causally efficacious and that we can grasp "real definitions" or real essences empirically.

4. Permissive naturalism's answer

Paradigmatic versions of moral naturalism endorse the thesis that moral properties are identical with or fully constituted by the natural properties on which they supervene (or a subset thereof). *Permissive naturalism* holds that for any moral property M, there is some natural property N on which M supervenes that is itself an infinite disjunction of natural properties, and M is identical with or fully constituted by N. *Restrictive naturalism* holds that for any moral property M, there is some natural property N on which M supervenes that is not itself a massive disjunction of natural properties, and M is identical with or fully constituted by N. Although permissive and restrictive naturalism differ in this way, they are identical in this important respect: they do not commit themselves to the further claim that fundamental moral concepts such as WRONG are natural in the sense that they, too, are such as to play an explanatory role within the natural sciences.[19]

Let us consider each version of naturalism to see how it might answer the projectability challenge. We begin with permissive naturalism, turning to its restrictive cousin in the next section.

---

[18] [citation omitted] provides defenses of both commitments.
[19] See Copp (2017) and, relatedly, Gibbard (2003, ch. 2). *Cp*. Cuneo and Shafer-Landau (2014, §5). Later, we touch upon this characterization of the natural.

Permissive naturalism embraces an 'intensionalist' theory of properties, according to which properties F and G are (metaphysically) necessarily equivalent if and only if either F is identical with or is fully constituted by G. Given the further assumption that there is a natural property N that consists in an infinite disjunction of natural properties on which wrongness supervenes, this view implies that wrongness is identical with or fully constituted by N. Since there is nothing unique about wrongness in this respect, permissive naturalism implies that every moral property is an infinitely disjunctive natural property on which it supervenes.[20]

Permissive naturalism has the virtues of being both clearly naturalistic and avoiding an overly restrictive account of what a property such as wrongness is. Moreover, in principle, the view has available several strategies for addressing the projectability challenge. Here we evaluate one that mirrors that endorsed by paradigm nonnaturalism. (In §7, we consider a very different type of strategy).

Recall that this strategy affirms

Grasp: In a range of cases, agents improve their reasonable mastery of moral concepts in virtue of having experiences as of their corresponding moral properties (or their instantiations).

It then claims that improvements in reasonably mastery of a concept due to such experience puts agents in a position to gain insight into important dimensions of its corresponding property, placing them in a position to project on the basis of experience and reflection. There is, however, a significant difference between how paradigm nonnaturalism and permissive naturalism

---

[20] This view is modeled on Jackson's (1998) position. For ease of explication, here and in much of what follows, we elide the phrase "or fully constituted."

characterize what is experienced: permissive naturalism maintains that in a case such as Spying your experience is as of an infinitely disjunctive property, while paradigm nonnaturalism does not.

The question to ask of permissive naturalism is why such experience would result in a change in your grasp of a moral concept, a change that (when all goes well) yields insight into dimensions of that property that in a case such as Spying would put you in a position to judge that your cousin's behavior is wrong. Here is one sort of answer: when agents experience the property *being wrong*, they experience it *as* an infinitely disjunctive natural property. This would provide an explanation of sorts of how you gain the appropriate insight: your experience of wrongness reveals its infinitely disjunctive character, enabling you to better grasp the concept WRONG, which enables you to discern that *spying for merely prurient reasons* is among wrongness's infinitely many disjuncts.

Let's grant that experience as of infinitely disjunctive properties is possible. As best we can tell, though, there is no plausibility whatsoever to the claim that, in a case such as Spying, the property *being wrong* is experienced as an infinitely disjunctive property. Nor do we have reason to take seriously the further claim that such experience reveals that *spying for merely prurient reasons* is among the infinite disjuncts of the property experienced. Such a proposal is on par with the implausible idea that, say, ordinary experience as of water reveals that the property *being water* has as one of its conjuncts the property *having two hydrogen molecules*.

A more promising approach for permissive naturalism would be to identify something about human cognition—or, alternatively—about wrongness itself that enables experience as of wrongness to yield insight into its character. That "something" might be what explains why

20

experience as of wrongness is not that of as a massively (and perhaps wildly heterogeneous) disjunctive property, but whose content has a recognizable unity or structure.

As an example of the former strategy, which focuses on human cognition, consider experience of color. According to some views of color vision, seeing something as (say) mauve consists in being sensitive to a great many discreet spectral profiles (i.e., ranges of energies of reflected or filtered light). There is an explanation for this phenomenon, known as 'metamerism': it is that normal human eyes contain three types of cones, each type of which is sensitive to a spectrum of light waves.[21] These cones enable us to correctly identify often quite strikingly disjunctive sets of spectral profiles as being instances of the same color. This mechanism (or something like it) might enable us to experience discreet spectral profiles as one color.[22] But, to our knowledge, there is no analogue to it in the moral domain. That is, there is no equivalent to a three-cone system that enables agents to experience infinite disjunctions of natural properties as the property *being wrong* and, so, improve their grasp of its corresponding concept.

As an example of the latter strategy, which focuses on the character of properties, consider a view according to which properties have modes of presentation (or 'guises'). Under this view, an infinitely disjunctive natural property might have the mode of presentation of *being wrong*. And direct experience as of this mode of presentation—so the proposal runs—could enable agents to better grasp its corresponding concept, thereby positioning them to project.

Philosophers put the expression 'mode of presentation' to different uses, sometimes using it is used to pick out concepts, sometimes properties. By all appearances, the proposal under consideration uses the expression to refer to *properties*. Are properties of this sort also highly disjunctive? If they are, then we are back to where we began; the proposal merely adds an

---

[21] Thornton (1973).
[22] We're grateful to an anonymous referee for pressing a concern along these lines.

epicycle that doesn't contribute to an explanation of how agents can grasp infinitely disjunctive properties in way that enhances their reasonable mastery of their corresponding concepts. If, however, these properties are not infinitely disjunctive, then we need an explanation as to why wrongness is an infinite disjunction of natural properties, as permissive naturalism maintains, rather than a mode of presentation, which is not highly disjunctive.

While we do not know whether there is any such explanation, neither option looks attractive. Suppose wrongness were identical with a mode of presentation. It follows that permissive naturalism is false, since this option implies that wrongness is not a highly disjunctive natural property, as permissive naturalism claims. Suppose wrongness were not a mode of presentation, but were merely to exemplify it. While available in principle given some further assumptions, this option highlights what might be our deepest reservation regarding the proposal we're considering: it leaves us in the dark as to why infinite disjunctions of properties have modes of presentation, let alone any particular mode of presentation at all. After all, what it is be a particular infinite disjunction of properties N would appear to have no implications regarding whether N has any mode of presentation, let alone some particular mode of presentation that reveals what property N is. The proposal under consideration appears not only fundamentally incomplete; it also doesn't point the way toward a satisfying answer to these questions.

We have considered three ways in which permissive naturalism might answer the projectability challenge. Although none looks promising to us, our aim is not to contend that permissive naturalism cannot answer the challenge. Rather, it is to draw attention to why its commitments make it difficult it to handle the challenge. Or, more exactly, our aim is to call attention to why, given its claim that moral properties are infinitely disjunctive, any such

explanation will probably not share the same form as the one presented in the last section in which Grasp and Insight play a central role.

5. Restrictive naturalism's answer

Restrictive naturalism earns its name because it claims that moral properties are natural but not massively disjunctive. In this way, the view thereby "restricts" what could count as a property such as wrongness. In the present context, the restriction is significant, since it is permissive naturalism's claim that moral properties are infinitely disjunctive that makes answering the projectability challenge particularly challenging.[23]

Yet restrictive naturalism faces challenges of its own, which we can understand in terms of how it might resolve a tension inherent to the view. On the one hand, restrictive naturalists might hold that wrongness is a paradigmatically natural property: one that is recognizably natural insofar as it is clearly suited to play an explanatory role in the natural sciences (or is clearly fully constituted by such properties).[24] We worry that this approach, which we'll call *reductive restrictive naturalism* (or simply 'reductive naturalism'), will lack the resources to ensure that a moral property such as wrongness has the correct type of normative-profile. And, so, the concern is that it will fail to provide sufficient reason to believe that those with reasonable mastery of the concept WRONG could project, accurately and justifiedly applying moral concepts to novel cases. On the other hand, restrictive naturalists might hold that wrongness is natural but not paradigmatically so: this might be because we lack the conceptual resources to identify some set of paradigmatically natural properties with which wrongness is identical (or that fully constitutes

---

[23] Since the same issues would arise if such views were to claim that such properties are not infinitely but massively disjunctive, we'll assume that restrictive naturalism also denies that moral properties are massively disjunctive.
[24] See our comment below on the rationale for working with this understanding of naturalism.

wrongness).[25] The concern we have about this approach, which we'll call *non-reductive restrictive naturalism* (or simply 'nonreductive naturalism'), is that while it might be able to ensure that a property such as wrongness has the correct sort of normative-profile, there is insufficient reason to believe it is a natural property. So, the worry is that the view will also provide insufficient reason for believing that agents could project, given its commitment to naturalism. Thus the tension: the commitment to naturalism pulls against the commitment to moral properties having the correct normative-profile and the commitment to moral properties having the correct normative-profile pulls against the commitment to naturalism.

Let us have a closer look at reductive naturalism. In principle, this view could furnish an answer to the projectability challenge of the right form. For if reductive naturalism were filled out in the right way, it might affirm that agents sometimes improve their reasonable mastery of moral concepts in virtue of having experience as of their corresponding moral properties, which positions them to project (per Grasp and Insight). The primary challenge facing this view is whether it can be filled out in the right way, incorporating an account of concepts and properties that has the right content.

To illustrate the worry, begin with the assumption that anyone with reasonably mastery of WRONG will grasp its central application conditionals, such as:

> necessarily, if some action falls under the concept KILLING OTHERS FOR SPORT, then WRONG applies to that action;

---

[25] *Cf.* Sturgeon (2006, 98-9).

necessarily, if some action falls under the concept DISSEMBLING FOR

REPUTATION'S SAKE, then WRONG applies to that action;

necessarily, if some action falls under the concept ABUSING ONE'S POSITION OF

POWER, then WRONG applies to that action; and so on.

Such a person will be in a position to accurately and justifiedly apply WRONG in novel cases only

if the actions specified in the central conditionals and those involved in novel cases are wrong.

But now consider a toy reductive naturalist view according to which wrongness is

"nothing over and above" the property *being such as to cause pain*. While this property is

recognizably natural, it is doubtful that its normative-profile includes the full range of actions

specified in the central cases, let alone the novel ones. Its normative-profile is too narrow (it fails

to include wrong actions other than those that cause pain) and undiscriminating (it fails to

distinguish between justified and unjustified inflictions of pain). Something similar appears true

of other paradigmatically natural properties that are somewhat more plausible candidates for

what wrongness is. Consider properties such as *failing to cause the most pleasure on the whole*

or perhaps *being disapproved of by an idealized agent*. These properties might have normative-

profiles that incorporate some of the central exemplification conditionals associated with

wrongness. But it is difficult to say what else they might include, as no one knows what actions

would fail to cause the most pleasure on the whole, or what an idealized agent would disapprove

of. In short: restrictive naturalism must hold that anyone with reasonable mastery of the concept

WRONG grasps its central application conditionals. But, if the view is true, there is insufficient

reason to believe that the normative-profile of wrongness would include the full range of central

exemplification conditionals, let alone those conditionals that pertain to novel cases. These views, then, also fail to provide sufficient reason to believe that agents with reasonable mastery of the concept WRONG could project.

Non-reductive naturalism is poised to do better. As indicated above, the leading thought behind such views is that properties such as wrongness are natural even if they are not paradigmatically so. Like paradigmatic non-naturalism, this view appears able to provide an answer to the projectability challenge of the right form, appealing to claims such as Grasp, Insight, and Isomorphism. But unlike its reductive counterpart, non-reductive naturalism can operate with an understanding of the natural capacious enough to guarantee that properties such as wrongness have the right sort of normative-profile, thereby ensuring that Isomorphism (or something like it) holds. In fact, thinkers such as Nicholas Sturgeon maintain that, in virtue of its being naturalistic, the view will draw from "our best moral theory," which presumably will affirm the central application and exemplification conditionals associated with wrongness.[26] If so, nonreductive naturalism also appears well-placed to provide an answer to the projectability challenge with the right content—indeed, one that looks indiscernible from that offered by paradigmatic nonnaturalism.

The challenge non-reductive naturalism faces is to explain why moral properties thus understood are genuinely natural. (We suspect that, given its open-ended character, this is probably *the* challenge the view faces from realist rivals.) We can bring the challenge into sharper focus by recalling that naturalism is a view about the *nature* of moral properties.[27] According to the gloss of naturalism with which we've been working—what Sturgeon calls the

---

[26] *Cf.* Sturgeon (1988, 241).
[27] Brink (1984, 121), Boyd (1988, 196, §4.3) and (2003, 527), and Copp (2007, 16, 26) are explicit about this. Boyd contends that these essences are known empirically. *Cp.*, also, McPherson (2015, 134ff). Here we gloss over nuances concerning what it is to belong to a thing's nature that in other contexts might matter.

"standard definition" employed in metaethics—a property's being natural consists in its being such as to play an explanatory role in the natural sciences.[28] Under the standard definition, then, naturalist views commit themselves to the claim that what it is to be a moral property is to be fitted to play (what we'll call) a *naturalistic explanatory role*.

Some naturalists, such as Sturgeon and Peter Railton, have argued that some moral properties play naturalistic explanatory roles.[29] If their arguments hit the mark, there is reason to believe that moral properties play causal (or nomic) roles in explaining various items of human behavior (perhaps your believing that an act is wrong) and the development of social dynamics over time (perhaps the social unrest in an unjust society). But to our knowledge, no naturalist has argued that *what it is to be* wrong is to be fitted to play such a role. If the characterization of naturalism offered above is correct, however, this is what must be defended by non-reductive naturalism. It is not enough to point out that, as a matter of fact, some such properties play naturalistic explanatory roles. This last claim, after all, is compatible with paradigm non-naturalism.

As with the concerns raised earlier regarding other versions of naturalism, this objection is intended not to be decisive but instead to identify a challenge that non-reductive naturalism must address if it is to explain how we can project, given its commitment to naturalism. Rather than anticipate what moves non-reductive naturalists might make in response, let us conclude this section with the following observation.

The projectability challenge has not been widely discussed. Although unfamiliar in important respects, it raises questions about nonnaturalism and naturalism that will be familiar from other contexts, ones that concern what these views say about the nature of moral reality.

---

[28] Or its instantiations being such as to play these roles; we elide this qualification in what follows. *Cf.* Sturgeon (2006, 92).

[29] Sturgeon (1988) and Railton (1986); cf. Brink (1984). *Cp.* Boyd (1988, §4.3).

Our aim in presenting the projectability challenge, then, is not to identify a set of issues that have hitherto entirely escaped the attention of metaethicists. It is rather to bring an unfamiliar problem to bear on some familiar issues, helping us to see why commitments regarding the nature of moral reality matter to satisfactorily answering questions in moral epistemology. The overarching thought is that, when commitments regarding the nature of moral reality help us to satisfactorily address these questions in moral epistemology, this can speak in favor of those ontological commitments. To the extent that these commitments make it difficult to satisfactory address these epistemic questions, this can put additional pressure on those ontological commitments. In this section, our contention has been that paradigmatic nonnaturalism's commitments regarding moral thought and reality enables the view to address one such question, while those of restrictive naturalism make it difficult for the view to do the same.

6. Conclusion

We have focused our attention on developing a non-naturalist response to the projectability challenge, arguing that paradigm nonnaturalism provides an explanation of the right form and includes the right content. We've also contended that we see no satisfactory naturalist response to the challenge. We recognize, however, that the argument we've offered incorporates a variety of controversial commitments. We conclude by canvassing three objections.

The first is that the argument overgeneralizes. For given the assumption that we can project in domains other than morality, it follows that we would have to be nonnaturalists about such things as numbers, the game of baseball, or autoimmune diseases. But that can't be right; these and other domains are populated by mere natural properties, relations, and facts. So even

without diagnosing what is wrong with the answer we offer to the projectability challenge, we can conclude that there must be something wrong with it.

We doubt that the argument we've offered overgeneralizes. It is true that, given the assumption that we can project in these other domains, our argument might imply that we should accept a broadly *essentialist* account of the concepts, properties, relations, and facts that constitute these domains. That is, given the assumption that we can project in these other domains, it might be that the best way to explain how those with reasonable mastery of the relevant concepts can project is that they grasp dimensions of the essence-profiles of the properties that correspond to these concepts.[30] But we deny that essentialism of this sort implies nonnaturalism. For it might be that a wide range of properties have a naturalistic essence-profile. Our view, then, is that whether we should be nonnaturalists with respect to the entities that constitute a given domain should be settled on a case-by-case basis.

The second objection points out that our challenge to naturalism hinges on controversial assumptions about the character of concepts and properties. Specifically, the challenge assumes that concepts must have a structure or unity of such a kind that makes sense of why they apply to their corresponding properties. Isomorphism is a vivid example of the assumption, as it implies that the application conditions that constitute the essence-profile of moral concepts are isomorphic with respect to the exemplification conditions that constitute the essence-profile of their corresponding properties.

Some prominent views of concepts reject this assumption. For example, a long-standing tradition has it that concepts are rules (or implicit conceptions) that enable us to think about the world in a variety of ways.[31] Under this view, to possess a concept is to grasp a rule, where the

---

[30] *Cp.* Korman (2014, 18).
[31] *Cf.* Peacocke (2008) and Ginsborg (n.d.).

character of the rule may differ very considerably from the structure or unity of its corresponding property.

A view such as this opens up possibilities for naturalism, particularly of the permissive variety. Suppose, for example, that the property *being wrong* were an infinite or massive disjunction of natural properties. And suppose an agent were presented with some small subset of this property. If the concept WRONG were a set of rules, then it could be that having reasonable mastery of it consists in being reliably disposed to apply these rules to all and only the relevant disjuncts of the property on a given occasion. Importantly for our purposes, it is not difficult to see how an agent would be in a position to project if this were correct; it would simply consist in having a satisfactory grasp of the relevant rules.

It is true that our discussion works with an understanding of concepts that is not neutral between competing accounts. There is, however, a rationale for this approach: we have assumed that a satisfactory account of concepts should not imply that concepts consist, or the possession of them consists, in agents having massive amounts of information (whether implicit or explicit) in the head, as the rule-based approach implies. In fact, bringing this assumption into the open helps us to appreciate virtues of the non-naturalist approach. This approach works with an understanding of concepts and reality according to which the two are integrated in such a way that reasonable mastery of moral concepts puts agents in a position to project. It does that, moreover, without assuming that being in such a position requires that agents must come to a novel situation with all the relevant information already in mind.[32]

The third and final objection is that our argument hinges on an overly optimistic estimation of our abilities to project, since it assumes that this is something that agents who have

---

[32] Burge (2003) and Smith (2015) develop closely related concerns about what we're calling the "rule-based" approach.

reasonable mastery of moral concepts can do in favorable conditions. But—the objection runs—this might be overly optimistic, since we might be rather poor at projecting in the moral domain. If this datum is overly optimistic, however, then there really isn't anything for metaethical theories to explain. So, it would be no objection to naturalism if it failed to answer the projectability challenge.

We concede that our argument hinges on an assumption that is moderately optimistic about our abilities to project in the moral domain. We find this assumption plausible. To deny it commits one to holding that we can form justified and accurate moral beliefs only in paradigm cases (or in conservative extensions of them). Moreover, the assumption can help explain how we make progress in the moral domain, and exploits parallels between the moral domain and other domains, including other normative domains. But we concede that we have not offered a full-fledged defense of this datum. Satisfactorily addressing skepticism of this variety would require further work.

References

Bealer, George. 1998. "A Theory of Concepts and Concepts Possession." *Philosophical Issues* 9: 261-301.

Boyd, Richard. 1988. "How to be a Moral Realist." In Geoffrey Sayre-McCord, ed., *Essays on Moral Realism*. Cornell University Press: 181 –228.

Boyd, Richard. 2003. "Finite Beings, Finite Goods: The Semantics, Metaphysics and Ethics of Naturalistic Consequentialism, Part I." *Philosophy and Phenomenological Research* LXVI: 505-53.

Brink, David. 1989. *Moral Realism and the Foundations of Ethics*. Cambridge University Press.

Burge, Tyler. 2003. "Concepts, Conceptions, Reflective Understanding: A Reply to Peacocke." In M. Hahn and B. Ramberg, eds. *Reflections and Replies: Essays on the Philosophy of Tyler Burge*. MIT Press.

Chalmers, David and Frank Jackson. 2001. "Conceptual Analysis and Reductive Explanation." *Philosophical Review* 110: 315-61.

Copp, David. 2007. *Morality in a Natural World*. Cambridge University Press.

Copp, David. 2017. "Are There Substantive Moral Conceptual Truths?" in Diego Machuca, ed. *Moral Skepticism*. Routledge: 91-114.

Cuneo, Terence. 2006. "Moral Facts as Configuring Causes." *Pacific Philosophical Quarterly* 87: 141-62.

Cuneo, Terence and Russ Shafer-Landau. 2014. "The Moral Fixed Points: New Directions for Moral Nonnaturalism." *Philosophical Studies* 171: 399-443.

Dancy, Jonathan. 2004. *Ethics without Principles*. Oxford University Press.

Gibbard, Allan. 2003. *Thinking How to Live*. Harvard University Press.

Ginsborg, Hannah. n.d. "Concepts and Rules: A Kantian Proposal."

Hurka, Thomas. 2014. *British Ethical Theorists from Sidgwick to Ewing*. Oxford University Press.

Jackson, Frank. 1998. *From Metaphysics to Ethics*. Oxford University Press.

Jackson, Frank, Phillip Pettit, and Michael Smith. 2000. "Ethical Particularism and Patterns." In Margaret Little, ed. *Moral Particularism*. Oxford University Press

Korman, Dan. 2006. "What Externalists Should Say about Dry Earth." *Journal of Philosophy* 103: 503-20.

Korman, Dan. 2014. "Debunking Perceptual Beliefs about Ordinary Objects." *Philosophers' Imprint*: 1-19.

McPherson, Tristram. 2015. "What's at Stake in Debates Among Normative Realists?" *Nous* 49: 123-46.

Oddie, Graham. 2005. *Value, Reality, and Desire*. Oxford University Press.

Peacocke, Christopher. 1992. *A Study of Concepts*. MIT Press.

Peacocke, Christopher. 2008. *Truly Understood*. Oxford University Press.

Railton, Peter. 1986. "Moral Realism." *Philosophical Review* 95: 163-207.

Railton, Peter. 2017. "Moral *Learning*: Conceptual Foundations and Normative Relevance." *Cognition* 167: 172-90.

Shafer-Landau, Russ. 2012. "Evolutionary Debunking, Moral Realism, and Moral Knowledge." *Journal of Ethics and Social Philosophy* 7: 1-37.

Smith, Michael. 1994. *The Moral Problem*. Blackwell.

Smith, Sheldon. 2015. "Incomplete Understanding of C oncepts: The Case of the Derivative." *Mind* 124: 1163-1199.

Sturgeon, Nicholas. 1988. "Moral Explanations." In Geoffrey Sayre-McCord, ed. *Essays on Moral Realism*. Cornell University Press: 229-55.

Sturgeon, Nicholas. 2006. "Moral Naturalism." In David Copp, ed., *The Oxford Handbook of Ethical Theory*. Oxford University Press: 91-121.

Thornton, William A. 1973. "Matching lights, Metamers, and Human Visual Response." *Journal of Color & Appearance* 2: 23-29.

van Roojen, Mark. 2015. *Metaethics: A Contemporary Introduction*. Routledge.

Wedgwood, Ralph. 2007. *The Nature of Normativity*. Oxford University Press.