

The Right Wrong-Makers

(Forthcoming in *Philosophy and Phenomenological Research*)

Richard Yetter Chappell*
University of Miami

July 24, 2020

Abstract

Right- and wrong-making features (“moral grounds”) are widely believed to play important normative roles, e.g. in morally apt or virtuous motivation. This paper argues that moral grounds have been systematically misidentified. Canonical statements of our moral theories tend to summarize, rather than directly state, the full range of moral grounds posited by the theory. Further work is required to “unpack” a theory’s criterion of rightness and identify the features that are of *ground-level* moral significance. As a result, it is not actually true that *maximizing value* is the fundamental right-making feature even for maximizing consequentialists. Focusing on the simple example of utilitarianism, I show how careful attention to the ground level can drastically influence how we think about our moral theories.

Stocker (1976, 454) famously lamented the “moral schizophrenia”, or disharmony “between one’s motives and one’s [normative] reasons,” that he associated with modern ethical theories. Our moral theories appear to furnish us

*Special thanks to an anonymous referee for *PPR*, who helped me to vastly improve the structure and organization of this paper. Thanks also to Anjan Chakravartty, Nicholas Cisneros, Roger Crisp, Abraham Graber, Johan Gustafsson, Nathan Robert Howard, Uri Leibowitz, Douglas Portmore, Mark van Roojen, Philip Stratton-Lake, Pekka Väyrynen, Allen Wood, Helen Yetter-Chappell, commenters at philosophyetc.net and philpeople.org, and audiences at Leeds, Miami, Reading, York, the 2018 OSU Dubrovnik Conference, and the 2019 Rocky Mountain Ethics Congress.

with highly abstract fundamental justifications—invoking the likes of aggregate utility, reasonable rejectability, universalizable maxims, or the balance of prima facie duties. Ordinary moral motivation, by contrast, often involves concern for particular, concrete individuals—and rightly so. This divergence between justification and apt motivation is all the more striking because many contemporary moral theorists explicitly endorse principles linking the two. Others (especially consequentialists)¹ have responded by disavowing this link, effectively embracing the charge of schizophrenic disharmony. But I think such disavowals are a mistake.

This paper offers a different kind of response to Stocker’s charge. We can reject the assumption that our moral theories furnish us with highly abstract fundamental justifications, normative reasons, or moral grounds. Our theories may advert to highly abstract properties in specifying their criteria for right action: that which fills in the blank in statements of the form, “An act is right iff ____.” But we need not take those canonical criteria to themselves *be* the theory’s fundamental moral grounds. Instead, I propose, we should interpret them as *summarizing* the full range of moral grounds posited by the theory. Highly abstract summary criteria are compatible with appropriately concrete and personal ground-level concerns. Harmony may thus be restored.

The central thesis of this paper is that the moral grounds (fundamental right- and wrong-making features) posited by a theory can be more specific than its general criteria for right or wrong action. What’s *criterial* for rightness need not be what most fundamentally *makes* an act right. This has

¹ Cf. Railton (1984)’s “sophisticated” consequentialism, or most recently, Sharadin (2019).

important normative implications, especially (but not only) for morally apt or fitting motivation. As a result, our choice of moral grounds can make an immense difference to the plausibility of our moral theories. Some powerful objections to consequentialism, for example, depend upon the assumption that *maximizing value* is the theory's posited right-making feature. Understanding why this assumption is false also makes clear why these objections fail. But it isn't only consequentialists who can benefit from the distinctions drawn in this paper. Moral theorists more generally would do well to ensure that their theories assign ground-level significance to particular individuals and their interests.

The structure of this paper is as follows: §1 explores the normative role of moral grounds, drawing connections both to apt motivation and to the phenomenon of *wronging* individuals. These connections raise what we might call *the Problem of Excessive Abstractness* for ethical theories. §2 explicates my central distinction between criterial and ground-level explanations of an act's moral status. This distinction is then put to work to provide theorists with a powerful, harmony-preserving solution to the Problem of Excessive Abstractness. Finally, §3 applies these lessons to dissolve 'motive' and 'wronging'-based objections to consequentialism from the recent literature.

1 The Normative Role of Moral Grounds

This section explores the normative roles of moral grounds, and how these give rise to the central problem that this paper seeks to address: the Problem of Excessive Abstractness.

1.1 *Apt Motivation*

Many philosophers now accept a link between virtuous motivation and the concrete grounds in virtue of which a good act has its positive moral status. For example, if donating to a certain charity is good *because it helps the global poor*, then it's natural to think that a donor ideally should be moved by just this factor (rather than, say, a desire to show off their wealth). Such moral grounds seem to be the proper objects of virtuous desires or morally worthy motivation. As [Smith \(1996, 182\)](#) writes, “We normally assume that [morally perfect] people are moved by the very features of their acts which make them right.”²

This paper elucidates a neglected ambiguity in this view of *Grounds as Motivators*, with important implications for our normative theorizing. To clear the way for this, let me briefly indicate my reasons for setting aside two alternative accounts of morally apt motivation: (1) ‘Moralism’, the view that agents should typically be motivated by explicitly moral thoughts, e.g. desiring to promote *goodness*, or to do one’s *duty*, as such; and (2) the ‘Global Consequentialist’ view that agents should have whatever motivational profile would be most useful, or bring about the best consequences ([Pettit and Smith 2000](#)).

² Related theses are often formulated in terms of ‘reasons’ rather than ‘right-makers’, e.g.: (1) [Arpaly \(2003, 72\)](#): “For an agent to be morally praiseworthy for doing the right thing is for her to have done the right thing for the relevant moral reasons—that is, the reasons for which she acts are identical to the reasons for which the action is right.” (2) [Markovits \(2010, 230\)](#)’s *Coincident Reasons Thesis*: “an action [is] morally worthy whenever the noninstrumental reasons for which it is performed coincide with the noninstrumental reasons that morally justify its performance.” (3) [Stratton-Lake \(2011, 372\)](#), drawing on [Korsgaard \(1989\)](#), offers a similar *Symmetry Thesis* on which “morally good people will tend to be motivated to do what they ought to do by the reasons why they ought to do those acts.”

The latter is arguably best understood not as a *competing* view at all, but rather as addressing a different question. An agent may ask, “What motivational profile would it be best or most desirable, from a moral perspective, for me to have?” I take this to be the sort of question driving Global Consequentialists. This is to ask about the morally recommended or fortunate motivational profile. By contrast, the question I’m interested in is importantly different—something more along the lines of, “What motivational profile is most morally *fitting* or apt, reflecting an *orientation toward the good*, and is on this basis perhaps *worthy* of praise or high esteem?”³

Note that there is no reason to expect the same answer to both questions, as an internal orientation towards the good may have bad extrinsic effects. For example, we may imagine that an evil demon threatens to destroy the world unless you acquire (and subsequently maintain) the very same vicious, morally contemptible motivations that drive the demon himself. He offers a magic pill that will induce this effect in you. As a good person, you care more about the world than about the purity of your own moral character, and so—quite virtuously—opt to take the pill and become vicious.⁴ Your *subsequent* motivational profile is, by design, morally contemptible. (We may suppose that you come to intrinsically desire to corrupt the virtuous, cause innocents to suffer, etc.) Nonetheless, it is highly morally fortunate

³ Note that to be *worthy* of praise is distinct from its being *good* to express praise, in the same way that for a proposition to warrant belief (based on the evidence) is distinct from its being good (perhaps for practical reasons) to assert.

⁴ This “virtuous viciousness” has the same general structure as Parfit (1984)’s “rational irrationality”. See also Smith (1996, n.4 on p.182): “a morally perfect utilitarian, in this ordinary sense, might none the less be morally required to cause himself to be less than morally perfect; that only by becoming less than morally perfect in the ordinary sense can he have the motivations his theory tells him he ought to have.”

or desirable—though you no longer care that this is so—because your new-found viciousness is causally responsible for saving the world from the evil demon’s threat. So the moral aptness of one’s motivations cannot be identified with their desirability or usefulness from a moral point of view. Since this paper concerns the former dimension of evaluation, we can put Global Consequentialism aside.⁵

There are two major reasons to prefer Grounds as Motivators over Moralism. The first appeals to intuitions about cases. [Stocker \(1976, 462\)](#) highlights how disturbing it would be to learn that a friend’s hospital visit was motivated by a sense of abstract duty rather than friendly concern. Moralistic motivation risks alienating agents from the grounded cares and concerns that we ordinarily think are central to the good life. As [Smith \(1994, 75\)](#) put it, “Good people care non-derivatively about honesty, the weal and woe of their children and friends, the well-being of their fellows, people getting what they deserve, justice, equality, and the like, not just one thing: doing what they believe to be right, where this is read *de dicto* and not *de re*.”⁶

The second reason is more theoretical, drawing on the conceptual equivalence between something’s being good (or desirable) and its being apt to desire ([Chappell 2012](#)). That is, it’s fitting to desire just those things that are good or desirable. But the (pure) moralist fails to have these desires. Instead of being moved by features that merit desire, they are moved by the

⁵ For more on how consequentialists should think about these two dimensions of evaluation, see [Chappell \(2012\)](#).

⁶ See also the authors referenced in note 2, who treat Moralism as a primary foil and argue against it at length—especially [Arpaly \(2003\)](#). Note that I take no stand here on Smith’s broader argument, particularly the question whether motivational externalists are committed to Moralism.

higher-order property of desirability, i.e., the property of possessing *other* features that merit desire. (This is akin to walking towards an arrow on a sign instead of following the direction in which the arrow points.) Once we see the structure for what it is, we can see that it would be more apt to be motivated by those meritorious features, the concrete good-makers, rather than by the higher-order property of possessing (other) meritorious features. But that’s just to favour Grounds as Motivators over Moralism.⁷

I hope that these brief introductory remarks suffice to assure the reader that Grounds as Motivators is (at least) a sufficiently well-motivated view to be worth exploring in greater depth. There has been some pushback against the idea that *any* degree of de dicto moral motivation is necessarily “fetishistic”—see, e.g., Svavarsdottir (1999)—but I do not rely upon any such extreme claim. I instead merely draw upon the moderate thesis that *de re* moral motivation is *a* central component of the good person’s motivational profile. (I think there’s a lot to be said for the hybrid view that it’s morally ideal to be motivated by the right-making features *in the recognition that they make the act right*, but it is not my task to defend this view here.)

⁷ For ease of exposition, I will sometimes slip between talk of features *of actions* and features *of the broader situation*. As Doug Portmore noted to me, we ordinarily think that what matters are things in the world (e.g., particular people) rather than mere features of actions (even the “right-making” features). But I don’t think that this actually undermines my assumption that an act’s right- and wrong-making features are moral grounds, because an act’s feature of *relating to the important things in certain ways* (e.g., the feature of causing harms or benefits to particular people) is obviously centrally relevant to there being good reason, or motivation, for performing *that act* in particular.

1.2 *The Problem of Excessive Abstractness*

We have thus far focused on the significance of moral grounds for determining morally apt motivation. This is just one of the important normative roles of moral grounds. [Stratton-Lake \(2011, 374\)](#) notes that they also “pick out whom we have wronged when we wrong someone.” That is, in order for a wrong action to have wronged a particular individual (giving them personal grounds for resentment), that individual must feature in the explanation of the act’s wrongness. Call this the *Wronged Appear in Wrong-makers* principle. For example, victims are plausibly wronged when *their* rights are violated, or when an agent fails in a duty that is owed *to them* in particular. In such a case, the violation of the victim is central to explaining why the act is wrong. Other kinds of wrong acts, by contrast, might be wrong merely due to their (expected) effect on the global welfare, in which case plausibly no one in particular is wronged by them.

We may further note that moral grounds appear, more generally, to be the focus of our moral theorizing. That is, traditional moral theories are typically in the business of attempting to identify those characteristics that *make* an act right ([Bales 1971](#)), and axiological theories similarly attempt to identify the features that make a state of affairs good. When utilitarians tell us that “an act is right iff it maximizes utility,” for example, it’s very natural to interpret this as saying that maximizing utility is what makes an act right, and failing to do so is what makes an act wrong. This natural interpretative step is based on the idea that *Theories State Grounds*.

If all this is correct, then we get some interesting results. For if you

accept that *Theories State Grounds*, and you further accept *Grounds as Motivators* or the *Wronged Appear in Wrong-makers* principle, then you're in a position to "read off" rather more from the statement of a moral theory than one might otherwise have thought possible. Specifically, you can read off implications for virtuous motivation or for when people are wronged.

Consider, for example, the following argument:⁸

1. *Grounds as Motivators*: A virtuous person is moved to ϕ by the very features of ϕ -ing that make it right for her to ϕ .
2. *Theories State Grounds*: A theory of the form "a subject's ϕ -ing is right iff her ϕ -ing has feature F" tells us that what makes it right for her to ϕ is that ϕ -ing has feature F.
3. According to consequentialism, a subject's ϕ -ing is right iff her ϕ -ing would maximize value.
4. Therefore, consequentialism implies that a virtuous person is moved to ϕ by the fact that her ϕ -ing would maximize value.

This seems a worrying conclusion: surely a virtuous person should care (non-derivatively) about particular persons, rather than anything so abstract as maximizing value!

Similar arguments could be presented against the familiar Kantian claim that *an action is wrong iff its associated maxim is not universalizable*.⁹ Since this statement makes no reference to particular individuals, it looks like the

⁸ Thanks to Doug Portmore for prompting me to make this explicit.

⁹ I take no stand here on whether this "familiar" Kantian claim is really true to Kant—cf. Wood (2017).

“real” victim of wrongdoing, for the Kantian who accepts *Wronged Appear in Wrong-makers*, is something like abstract Reason rather than particular persons. And we can use *Grounds as Motivators* to infer that a consistent or “virtuous” Kantian would fundamentally care only about universalizability, rather than caring non-derivatively about particular persons and their distinctive interests.

So, some prominent moral theories appear to be in trouble here. To label the issue, we might say that they suffer from an apparent *Problem of Excessive Abstractness*. But appearances can be misleading. In what follows, I argue that Theories State Grounds is subtly mistaken, in a way that undermines the troubling arguments above.

My strategy here diverges significantly from the standard consequentialist move of divorcing decision procedures from the criteria for right and wrong action (Bales 1971; Railton 1984).¹⁰ Whereas I reject premise (2) of the above argument, Railton and others implicitly reject (1). They address the Problem of Excessive Abstractness through a kind of indirection, setting aside close examination of the theory’s ultimate justifiers to instead invite whatever contingent motivations would best serve to promote the good. This indirection leaves those consequentialists vulnerable to Stocker (1976)’s charge of disharmony between their motivating and normative reasons.¹¹ I offer a

¹⁰ Of course, I’m happy to grant the basic distinction. I just think that the two can be reconciled, or successfully bridged, by principles like Grounds as Motivators. Specifically, such principles allow us to relate criteria of rightness to *fitting* motivations. It remains an independent empirical question which motivations would be most instrumentally *good* to possess in any given situation. So the deeper distinction here is not between criteria of rightness and decision procedures, but between what’s fitting (in principle) and what’s value-promoting (in practice).

¹¹ Related objections to Railton’s sophisticated consequentialism are developed at length in Chappell (2019a, sec. 3). The basic worry is the same that we raised against Global

more harmonious solution, addressing the Problem of Excessive Abstractness head-on by arguing that our theories' ultimate justifiers need not be so abstract in the first place.

2 Criterial and Ground-Level Explanations

This section argues that canonical statements of our moral theories may serve to summarize, rather than directly present, the full range of moral grounds posited by the theory in question. As a result, our moral theories are, as they stand, surprisingly incomplete in presentation. More work must be done to flesh out the full story about what right- and wrong-making features they posit.

2.1 *Introducing the Distinction*

A simple way to bring out the challenge to Theories State Grounds is to consider how it misfires when applied to the basic statement of Rossian deontology: *an act is wrong iff it violates the weighted balance of one's prima facie duties*.¹² It would seem mistaken to hold that the wrong-making feature

Consequentialism earlier: by substituting value-promoting motivations in place of apt ones, these consequentialists risk changing the subject. Of course, it's true that the value-promoting motivations are the best ones to have, but that simply *isn't the same thing* as their being fitting, virtuous, or what we mean when talking about acting with "moral worth" or for the "right reasons". We may happily grant that promoting value *matters more*, in practice, than the latter sort of normative status. Even so, as moral *theorists* we should want to have an accurate account of this other kind of normative status. And that requires us to look beyond mere value-promotion, even for consequentialists.

¹²Ross (1930/2007, 41): "We may try to state first what (if anything) is the universal nature of *all* acts that are right. ... [R]ight acts can be distinguished from wrong acts only as being those which, of all those possible for the agent in the circumstances, have the greatest balance of *prima facie* rightness, in those respects in which they are *prima facie* right, over their *prima facie* wrongness, in those respects in which they are *prima facie* wrong." Olsen (2014) argues on the basis of the quoted passage that Ross should be

of one's maleficent act, according to Rossianism, could only be the abstract fact that it violated the weighted balance of one's prima facie duties. We normally recognize that, for the Rossian, particular prima facie duties (read: *pro tanto* moral reasons) contribute non-derivatively to the explanation of what we ought to do. As Stratton-Lake (2011, 366) explains, the weighted balance serves an important summary role in specifying "what is recommended", but it's the particular things being balanced—the various moral reasons, or prima facie duties, themselves—that do the real work, as "what recommends."

This serves to highlight a general distinction between what we may call the *critical* and *ground-level* explanations of why some moral fact obtains. A *critical* explanation appeals to the (explanatorily relevant)¹³ necessary and sufficient conditions for the truth of some moral claim, i.e., the conditions that appear in place of the 'X' in theoretical accounts of the form, "An act is right if and only if X." If I randomly kick Joe in the shins, the wrongness of my act can be explained *critically* by the fact that my act has the general property Y, which all wrong actions share. (Maybe Y is the property of failing to maximize value, or maybe it is the property of violating the weighted balance of prima facie duties.)

A *ground-level* explanation, by contrast, appeals to the particular features

understood as a generalist rather than a particularist. But even if one doubts this—say on the grounds that we lack any sort of general formula for *how to weigh* the different prima facie duties against each other—it's hard to see why this should make any difference to my subsequent arguments. Whether individuals are ever wronged, for example, cannot plausibly depend upon imprecisely specifying the weightings of different prima facie duties.¹³ There might be other criteria that are necessary and sufficient due to co-varying with the genuinely explanatory criteria. For example, a theist who rejects Divine Command Theory for Euthyphro-style reasons might regard the property of *being approved of by God* as necessarily co-extensive with *maximizing utility* or whatever they take the independent criterion of rightness to be.

of the act which ground its having the moral status that it does, or that *make* it right or wrong. So, for example, the ground-level explanation of my action’s wrongness may consist in the fact that the act (gratuitously) harmed Joe. My harming Joe is what makes it the case that my action meets the more general criteria for wrongness, and is hence the fundamental *wrong-making feature* of my action.¹⁴

Both criterial and ground-level explanations can serve as legitimate answers to the question, “Why was that action wrong?” If the asker is wondering what makes actions in general wrong (when they are), and so hopes to learn what *generalizable property* this wrong action shares with all other wrong actions, then they are asking for the criterial explanation: “Because it failed to maximize value.” On the other hand, if they are more interested in the specifics of this particular action, how it came to meet these more general criteria, or what the basic morally significant features of the situation were, then they are asking for the ground-level explanation: “Because it (gratuitously) harmed Joe.”¹⁵

We need this distinction to make sense of the conceptual possibility of various wrong acts being wrong *for different reasons*. All wrong acts share the property criterial for wrongness (by definition). So if the criterial property was the only eligible wrong-making property, it would seem to follow that

¹⁴ One might allow that there is a (derivative) sense in which the general criteria also involve “wrong-makers”, since it is cogent to explain an act’s wrongness by appeal to its possession of the criterial property. It simply isn’t the most basic or *relevant* explanation for our purposes, as the fundamental grounds are what play the important normative roles identified in section 1. Unless otherwise specified, in this paper when I speak of “right-” and “wrong-making features”, I mean the ground-level ones.

¹⁵ I focus on a particular (token) action to make the distinction vivid, but the ground-level explanation could also answer the question of what makes acts *of that (suitably) specific type* wrong. Thanks to an anonymous referee for prompting me to clarify this.

all wrong acts were wrong for the same reason. Anyone disinclined to accept the latter conclusion is thereby committed to holding that the ground-level explanation of an act's wrongness may be more specific or fine-grained than the general criteria for wrongness. That way, two acts that equally meet the general criteria for wrongness may nonetheless differ in what it is that *makes* them wrong. We will see more concrete examples of this in the discussion to follow.

2.2 *Applying the Distinction*

It's not a trivial matter to discern which normatively significant properties have ground-level normative significance. Note that the normative ground-level need not be located at the "lowest" or most basic level of descriptive reality (fundamental microphysics or the like). Normative grounds must have the right kind of (non-derivative) normative significance, and a "higher-level" property may well be more normatively significant than a descriptively "lower-level" one. Consider a Divine Command Theorist who believes that God contingently commands us to relieve suffering. While the property of relieving suffering plays an important normative role on their view, its significance is entirely derivative. As a result, we do better to identify the property of *being commanded by God* as their ground-level right-maker, or what does the normative heavy-lifting. After all, the subjectively fitting moral motivation for such an agent would seemingly be to have a basic desire to do as God commands, and a merely derivative or instrumental desire to relieve

suffering.¹⁶

This clarification brings out that it is a substantive normative question what features have ground-level moral significance, or feature in ground-level normative explanations. My view is that the normatively correct level of specificity involves features like *harming Joe* (to a certain extent), rather than either more specific features like *hurting Joe's left shin* or more general features such as *harming a sentient being*.

One way to test this is to ask whether it makes a normative difference how the feature is realized. It makes no normative difference whether you hurt Joe's left shin or his right one, all else equal: either token act would be wrong in just the same way—which is to say, they share the same wrong-making feature (and should not be distinguished at the normative ground-level). But it does make a normative difference whether you hurt Joe or Jane. Those two harmful acts are wrong for importantly different reasons, we're inclined to think.

Our normative verdicts thus commit us to the view that *harming Joe* and *harming Jane* are distinct wrong-makers.¹⁷ That's why it would seem

¹⁶ At least, this is so on the simple 'obedience' interpretation of Divine Command Theory (DCT) that I use here for illustration. Perhaps we could imagine a more sophisticated view—a kind of 'virtue-ethical' interpretation of DCT—on which it's instead fitting to have concern for *just those things* that God himself has concern for. The striking difference between these two forms of DCT nicely illustrates the theoretical significance of the normative ground-level, and how it can vary even while holding fixed the general criteria for rightness.

¹⁷ To translate into the idiom of Chappell (2015), tradeoffs between distinct moral grounds call for *ambivalence*, whereas truly normatively equivalent outcomes are fungible such that any tradeoffs between them warrant outright *indifference*. It would be very odd to feel terribly torn about whether Joe is to suffer a left-shin pain or a right-shin pain. That's a difference that clearly doesn't matter. But it would be a deep moral error to similarly disregard a tradeoff between Joe's suffering and Jane's. We should feel torn in the latter case, because the separateness of persons means that this is a difference that matters.

a mistake to attribute ground-level significance to a more general feature like *harming a sentient being*. To explain the normative distinctness of the two wrong acts—the sense in which they are wrong for importantly different reasons—we must appeal to features of the situation that differ between the two cases. But our cases of harming Joe vs. harming Jane may be alike in all (normatively relevant) respects except for the identity of the victim. So the normative ground-level must advert to particular persons: whether it is Joe we hurt, or Jane. (Not that either person matters *more* than the other, of course. The claim is just that they each matter, separately.)

Such specificity is also required to solve the Problem of Excessive Abstractness, if we accept the normative roles for moral grounds set out in §1. The “Problem” ultimately stems from our implicit commitment to avoiding excessively abstract normative grounds, together with the presumption that our moral theories cannot meet this commitment. I’ve now explained in a different way why our ground-level explanations must be specific enough to distinguish between persons. It remains to be seen whether our moral theories can accommodate this need.

To answer this question, we need to know what constraints are imposed by a theory’s criterial explanations, and so what ground-level explanations remain compatible with any given theory’s criterial statements. To address this, we may clarify the relationship between the two as follows: Criterial explanations serve to (non-trivially) specify the general conditions under which any more particular ground-level explanation will obtain.¹⁸ My (gra-

¹⁸I add “non-trivially” to exclude “because it was wrong” from counting as a criterial explanation of an act’s wrongness. I needn’t take a stand here on whether infinite disjunctions (i.e., of every possible particular wrong action across all possible worlds) should

tuitously) harming Joe is the specific morally relevant feature in virtue of which my action meets the general conditions for wrongness. These general conditions for wrongness might have been met in different ways, say if my kick had hurt Jane instead of Joe.¹⁹ There would then be a slightly different wrong-making feature, or ground-level explanation of how my action came to satisfy the general criteria for wrongness.²⁰

This brings out that our ground-level wrong-makers must (at the very least) contribute to the satisfaction of our general criteria for wrongness, and so potentially explain why, in their specific context, the criteria were ultimately met. Of course, the explanation here should be constitutive rather than causal: if an evildoer saves one person as a means to killing five, his saving the one is not itself a wrong-making feature, though it causally explains the obtaining of other features—the killings of the five—that are wrong-

similarly be disqualified. This further question addresses whether particularists—who deny that any finite such generalities are available—are committed to denying that criterial explanations are possible at all, or just that they aren't useful, finitely specifiable, etc. The former option probably makes more sense, however, as it seems that the putative “criterial explanations” in the latter case would be incapable of doing any real explanatory work. Put another way, “Because it was one of the following maximally specific possible actions: A or B or C or . . . *ad infinitum*,” does not seem much of an answer to the question “Why was this act wrong?”

¹⁹That is to commit a different token wrong of the same general type. Of course, on pluralist accounts, the criteria for wrongness might also be satisfied in a different *type* of way—say, if I violated a prima facie duty of fidelity rather than of non-maleficence.

²⁰This distinction between the criterial and the ground-level is not unique to ethics. One could similarly distinguish between the general criteria for being a sibling, and the particular individuals—my brothers—who make it the case that I am a sibling. But it's worth reiterating that the sort of grounding this paper is concerned with is distinctively *normative* rather than *metaphysical*—cf. [Fine \(2012\)](#). So I don't take my ‘ground-level explanations’ or right-making features to settle the metaphysical nature of rightness, for example. Rather, I mean to leave open that one could be a non-naturalist about normative properties, whilst allowing that natural properties serve as right- or wrong-making features (or feature in ground-level normative explanations). For more on what non-naturalists should say about the relation between rightness and right-making features, see [Chappell \(2019b\)](#).

making. Crucially, harms to individual persons are going to constitutively contribute to meeting the criteria for wrongness on any plausible moral theory. I'll illustrate with two.

For Rossians, an agent's harming Joe may constitutively explain why they failed in their prima facie duty of non-maleficence, which may in turn be the decisive factor for explaining why they failed to fulfil the weighted balance of their prima facie duties in that situation. For consequentialists, the connection is even more straightforward: since Joe's welfare has intrinsic value, harming Joe directly contributes to (and, given appropriate background conditions, may constitutively explain)²¹ one's failure to maximize value.

These theories are thus eligible to treat *harming Joe* as a ground-level wrong-maker, and thereby absolve themselves of the Problem of Excessive Abstractness. The resulting *well-grounded* theories are straightforwardly superior to extensionally equivalent variants that settle for excessively abstract moral grounds, for all the reasons explained above (and further developed below, in §3). So interpretative charity recommends that we regard the well-grounded variant of each theory as the default or canonical version of the view.

2.3 Revisiting Grounds as Motivators

We've now established the central thesis of this paper: a theory's moral grounds may be much more specific than its criteria for rightness and wrongness, allowing for a direct solution to the Problem of Excessive Abstractness

²¹ §2.4, below, explains why we don't need to explicitly build all those background conditions into the official "wrong-making feature."

that preserves harmony between our normative and (fitting) motivating reasons. But it's worth expanding upon how our distinction between the criterial and ground levels can further illuminate our understanding of virtuous or apt motivation.

Recall our observation (from §1) that the properly conscientious agent's concern for "morality" should be (at least in large part) *de re* rather than purely *de dicto* in nature. That is, rather than just caring about "morality" or "rightness" in the abstract, the morally conscientious agent cares about *the things* that are of moral significance or the right-making features.²² It would seem perverse, after all, to neglect the concrete things that matter in favour of the abstract property of *mattering* (especially if one's conception of morality turns out to be terribly misguided: just think of the zealous enforcer of unjust laws, e.g., Javert in *Les Misérables*).

In a similar fashion, I propose, it would be perverse to neglect what is of ground-level moral significance in favour of what is merely criterial—i.e., a mere general guarantee that *some or other* feature of ground-level moral significance obtains. Such a *criterially motivated* agent or 'Criterial Moralists' has taken one step towards concreteness over the bare Moralists that we began with. The Criterial Moralists have more fleshed out moral motivations, in line with the descriptive substance of their criterion of rightness (maximizing happiness, or whatever it may be). But they have further yet to go. Our general criteria of rightness can be further unpacked into particular morally

²² As per the terminological clarification in note 7, these are not the same thing. So it would be more precise to say that the conscientious agent cares about the things that are of moral significance, and so is moved to act by an act's (right-making feature of) being suitably related to those morally significant things.

significant considerations, or ground-level reasons, and it is these that really ought to move us.

To illustrate concretely: It is perverse for a wannabe Rossian to care only about the balance of prima facie duties, rather than caring about avoiding harm to Jane, and maintaining fidelity to Sally, and so forth. Such an agent is criticizable on internal, Rossian grounds, for failing to care appropriately about what (according to the most plausible interpretation of Rossianism) has ground-level moral significance, namely the concrete *contents* of their prima facie duties.

Analogous criticisms apply to the utilitarian who cares only about aggregate utility and not particular people. Any plausible form of utilitarianism must be *token-pluralistic* in its axiology, ascribing intrinsic value—and hence ground-level significance—to each distinct individual’s welfare (Chappell 2015). So an agent who fails to care about each particular person’s welfare is thus failing to care about what, on this view, has ground-level moral significance. The imagined perverse agent is thus not accurately exemplifying the (most plausible) utilitarian perspective after all.

The distinction between ground-level and criterial explanations thus explains how it is that the putative Problem of Excessive Abstractness misfires. *Theories State Grounds* misled us into thinking that a theory’s criteria for rightness were also its fundamental right-making characteristics. We now find that greater care is required to correctly identify (e.g.) utilitarianism’s moral grounds. But once we do, the view can be combined with *Grounds as Motivators* to yield plausible verdicts for apt motivation after all. That, again, is the central upshot of this paper: a theory’s moral grounds may

well be more specific than its general criteria of rightness, in a way that significantly enhances the theory's attractiveness.

Careful attention to the ground level can help us to better interpret other moral theories too. We can now see that even Kantians should probably not be motivated primarily by anything so abstract as the Formula of Universal Law, and Contractualists can aptly have more concrete moral concerns than whether they are conforming to principles that no-one could reasonably reject. In all cases, the abstract criteria of rightness offered by our moral theories can be fleshed out to yield more concrete ground-level right-makers, and it is these that (most plausibly) should typically motivate virtuous or morally conscientious agents.

This suggests an important emendation to [Markovits \(2010, 230\)](#)'s *Coincident Reasons Thesis* that "an action [is] morally worthy whenever the noninstrumental reasons for which it is performed coincide with the non-instrumental reasons that morally justify its performance." Without our distinction between criterial and ground-level justifications, the Coincident Reasons Thesis risks being too permissive, bestowing (full) moral worth upon acts done from excessively abstract moral motivations. Whilst we might grant some moral worth to such abstractly motivated actions (as being more morally worthy than purely self-interested acts, for example), it seems that acts are typically *most* morally worthy when the agent's motivations coincide with the *ground-level* justifiers. Something has gone wrong with the Criterial Moralist who cares more about abstract principles than concrete people, after all.

2.4 *Further Clarifications: Normative Grounds and Decisive Reasons*

One may object to my identification of *harming Joe* as a possible wrong-making feature on the grounds that this very feature could also be present in an act that was permissible, say if harming Joe was necessary to prevent far greater harms to Jane. Does this imply that the *real* wrong-maker for the act must instead be some more general comparative property, e.g. *harming Joe when an alternative involving less net harm was available*? I think not.

One option for addressing this concern would be to distinguish full vs. partial grounds, and suggest that *harming Joe* is an instance of the latter. Perhaps the full ground-level explanation of the act's wrongness depends upon all sorts of details of the situation, including facts about the various different options available to the agent and how each of them compare to the act of harming Joe. Still, the response goes, the harm to Joe is at least a *partial* ground of the act's wrongness, and perhaps that suffices to explain why a virtuous agent would be motivated by this consideration.

My main concern with this response is that it doesn't seem to give sufficient *centrality* to Joe's suffering. If it is just one partial ground amongst many, could a virtuous agent just as well be motivated by concern for some other partial ground instead? And wouldn't it be *most* appropriate to be motivated by the *full* grounding explanation, rather than just a part of it? Perhaps it wouldn't be terrible to hold that one should be motivated *both* by concern for Joe and by a more general concern for the criterial moral principles. But I would rather not be committed to this, and I'm especially wary of any suggestion that the latter concern should be more fundamental

than one's (non-comparative) concern for Joe as an individual.

A better response, I think, harks back to the illuminating structure of Rossian "prima facie duties" (really: *pro tanto* reasons). The prima facie duty of non-maleficence means that we have a pro tanto reason not to harm Joe. If this is not outweighed by any competing reasons, then harming Joe will qualify as all-things-considered wrong: the prima facie duty becomes a final duty. In such a case, what ground-level explanation should the Rossian give of the act's wrongness? One might, as above, appeal to both *harming Joe* and *the absence of any other reasons*, which together suffice for wrongness. But that response neglects the very different roles that the two considerations are playing in the overall normative explanation. The harm to Joe constitutes a normative reason against so acting. The absence fact is not itself a normative reason at all, but instead tells us that there aren't any further reasons we need to consider. As such, the absence fact serves as an *enabler* for turning a pro tanto wrong-maker into an all-things-considered one.

This brings out an important clarification to the notion of *Grounds as Motivators*. The claim is not that virtuous agents should be motivated by everything that plays some metaphysical role in determining the moral status of an act. For, as we've seen, that could include higher-order facts *about* reasons that are not themselves normative reasons at all. We should instead insist that the relevant kinds of "moral grounds" are just the (decisive) ground-level *normative reasons*. It's a familiar fact that *whether some normative reason is decisive* is a contingent matter that depends upon how it compares to the other reasons in the situation. But this in no way implies that the only "real" normative reason is instead some more abstract com-

parative fact, or indeed that the latter is any kind of reason at all.

(One might suggest, in light of all this, that the terminology of ‘right-’ and ‘wrong-makers’ is needlessly confusing or perhaps even misleading. It is, perhaps, an unfortunate consequence of the way the literature developed in response to Moralism that we have this terminology. I have no objection to jettisoning this terminology and restating all the relevant theses directly in terms of normative reasons—cf. note 2. Nothing of substance hangs on this terminological choice. But if we do use the terminology of grounding to talk about virtuous motivation, it is important to charitably interpret it in the way I have suggested here. Otherwise the relevant normative theses have no chance at all of being true.)

3 The Implications for Consequentialism

The previous sections set out some general lessons for our normative theorizing. This final section seeks to illustrate them by stepping through, in greater detail, how consequentialists (especially utilitarians) can apply these lessons to undermine ‘motive’ and ‘wronging’-based objections from the recent literature. Along the way, we will clarify how to determine what’s of ground-level significance for consequentialists.

3.1 The Motive Objection

Section 1.2 sketched the motive objection to consequentialism, offering it as an example of the Problem of Excessive Abstractness. We may now flesh out this specific objection with reference to [Stratton-Lake \(2011, 380\)](#)’s presenta-

tion. He claims that, according to consequentialism, “[I]t is the fact that [an] act has produced the best outcome, not the fact that it has produced a state in which [various particular values are promoted], that makes it the right act to do.” When combined with Grounds as Motivators (or Stratton-Lake’s “symmetry thesis”, from p. 372, “that morally good people will tend to be motivated to do what they ought to do by the reasons why they ought to do those acts,”) we get the impalatable implication that, were consequentialism correct, morally good people would tend to be motivated by thoughts of producing the best outcome.

While that might seem a reasonable enough motivation for certain (e.g., philanthropic) acts, it seems implausibly impersonal and alienated as an account of everyday caring acts. To channel Williams (1981, 18), when a man saves his wife, for example, “it might have been hoped... that his motivating thought, fully spelled out, would be the thought that it was his wife, not that it was his wife and that in situations of this kind it is [not suboptimal] to save one’s wife.”²³

Crucially, this objection depends upon the assumption that the right-maker for consequentialism is “produc[ing] the best outcome.” No argument is offered in support of this claim; Stratton-Lake treats it as obvious, presumably on the basis of an implicit commitment to the idea that Theories State Grounds. But once we distinguish the criterial from the ground level, as we did in section 2, this assumption can no longer be taken for granted. Theories state *criteria*, which may need further unpacking to reveal their ul-

²³ The original Williams quote seemed to assume Moralism, attributing to the agent gratuitous thoughts about *permissibility*.

timiate grounds. If the latter are sufficiently concrete and particular, so can apt motivation be.

How, then, are we to identify a theory's ground-level right-makers? The general answer, I think, involves looking to the particular morally significant components (or normative reasons) that explain how the general criteria for rightness come to be satisfied in any given case. For act consequentialism, the general criterion is *producing the best outcome*, and the components that contribute to the satisfaction of this criterion are *particular token values* (specified in the theory's axiology). Consequentialists should presumably hold that agents have pro tanto reason to ϕ just when ϕ -ing promotes some token value. So their normative reasons for action are given by whatever (token) values are promotable in their particular circumstances. Ideally, these reasons then become the agent's motivating reasons, insofar as they are (informed and) virtuous. So, if we believe that virtuous agents would be moved by concern for particular individuals (and not just the general good), consequentialists can accommodate this by adopting a token-pluralistic axiology that specifies each individual's welfare as a separate basic good. The motive objection to consequentialism thus fails.²⁴

²⁴That is, it fails as an objection to consequentialism *in general*; specific axiologies may still be subject to objections based on their implications for moral motivation. For example, Stratton-Lake's objection was directed at an "expanding the good" strategy which sought (e.g.) to accommodate common intuitions about partiality by attributing intrinsic value to *sons helping their fathers*. But even on my analysis, this putative value doesn't seem like it yields the desired motivations. We presumably want a son to be motivated to help *his* father, not just to promote states of affairs in which *sons help their fathers*—as though enabling two *other* sons to help their respective fathers would be preferable to saving his own. So I agree that this variant of consequentialism does not seem a promising way to accommodate common-sense partiality. (For that, you need agent-relative values.)

3.2 *Wronging Individuals*

Can utilitarians account for the wronging of individuals? Nelson (2015, 717) claims not: “According to utilitarianism, moral offenses are offenses against global utility, right reason or the totality of sentient beings, but never against individual victims.” Stratton-Lake (2011, 382) similarly claims, “the fact that [an agent] has failed to bring about the best state of affairs leaves the fact that his [victim] has been wronged completely unexplained. If anyone is wronged here, it would be the world which has not been made as good as it could have been. But that, of course, makes no sense.”

In either case, the underlying reasoning seems to depend upon something like the Wronged Appear in Wrong-makers principle, together with the assumption that utilitarianism’s wrong-makers concern “global utility” and never “individual victims”. Now that we’ve seen our way clear of the latter assumption, the objection seems on shaky ground.

Utilitarians may, for example, hold that individuals are wronged by agents who fail to give their interests due weight. Token-plural utilitarianism ascribes equal ground-level significance to each individual whose interests are at stake in the situation. Neglect of such interests may play a central role in explaining why some harmful act was wrong. In such a case, the agent’s lack of due concern for their victim constitutes a targeted form of moral disrespect, directly analogous to (on the Rossian view) disregarding a prima facie duty of non-maleficence.

To illustrate, consider three different ways that I might sub-optimally (hence wrongly) harm you in order to benefit some other individual. (1) The

benefit to the other person may not even outweigh the harm done to you; in that case I have distinctively wronged you by acting in a way that fails to give sufficient weight (equal consideration) to your interests. (2) Suppose instead the benefit does outweigh the harm done to you, but I could have benefited some third party even more. Then you have no special complaint; I have adequately taken your interests into account (or at least acted in a way that's compatible with doing so), properly recognizing that they are outweighed by the greater benefits I can (and do) give others. But I have now distinctively wronged the third party whose interests have been unduly neglected. (3) Suppose the extra "benefit" I forego was instead to a merely possible person: I could have, but failed to, bring into existence a whole new life that would have been a very happy, flourishing life. Then no-one in particular is wronged. This is an instance of a *purely* impersonal wrong. As the first two scenarios demonstrate, the utilitarian can easily account for particular persons being wronged or offended against in appropriate cases.²⁵

²⁵ Some cases are less clear (but in a way that is not special to utilitarianism). Suppose I harm *both* V and W in order to give a benefit to some third party that is greater than either harm in isolation, but much less than both harms combined. Our account of *Wronging as Undervaluing* suggests that we do not yet have enough information to know who my action wrongs. Absent any basis for differentiation, it's most natural to assume that I undervalue (and hence wrong) both V and W. But suppose that isn't the case. Suppose I fully respect V, and wouldn't impose a larger harm on her for a smaller benefit to anyone else; it's just that I don't give W's interests any weight whatsoever, and so (wrongly) consider the combined harms to V + W to be of no greater importance than the harm to V alone. In such a case, I'm inclined to think that V has no reason to resent me for acting as I do (whereas W certainly does), so only W is wronged.

Note that Rossian deontologists face analogous cases where a single prima facie duty of non-maleficence would not suffice to make an act wrong, but two together do, and the agent acts wrongly because they ignore (or underweight) just one of these two prima facie duties. Again, I'm inclined to think that the correct answer for the Rossian is to hold that only the subject of the neglected prima facie duty is wronged. But anyone who prefers a different solution in the Rossian case (e.g., holding that both are wronged even though it would have been permissible to harm the one had it not been for the second) will presumably find a corresponding response available to the utilitarian.

Conclusion

The distinction between *critical* and *ground-level* features undermines the assumption that Theories State Grounds, requiring us to rethink how we identify a theory's right-making features and other moral grounds. We may still retain the background thought that specifying moral grounds is a central task for our moral theorizing. It is simply one that goes beyond the more familiar task of specifying the general criteria for rightness. Moral theorists thus have more work to do than previously appreciated, as a theory remains incomplete if it merely provides us with a summary criterion without specifying which features are of ground-level significance. We've seen that it makes a significant difference to the plausibility of utilitarianism, for example, whether aggregate welfare or individual welfare occupies the ground level. By adopting the latter view, utilitarians can avoid the Problem of Excessive Abstractness, securing plausible implications for virtuous motivation (while accepting Grounds as Motivators) and accounting for the wronging of individuals (while accepting Wronged Appear in Wrong-makers). Developing plausible specifications of the ground level for Kantian, contractualist, and other moral theories is a task for future research.

References

- Arpaly, Nomy. 2003. *Unprincipled Virtue: An Inquiry Into Moral Agency*. Oxford University Press.
- Bales, R. E. 1971. “Act-utilitarianism: account of right-making characteristics or decision-making procedures?” *American Philosophical Quarterly* 8:257–65.
- Chappell, Richard Yetter. 2012. “Fittingness: The Sole Normative Primitive.” *Philosophical Quarterly* 62:684–704.
- . 2015. “Value Receptacles.” *Noûs* 49:322–332.
- . 2019a. “Fittingness Objections to Consequentialism.” In Christian Seidel (ed.), *Consequentialism: New Directions, New Problems?* Oxford University Press.
- . 2019b. “Why Care About Non-Natural Reasons?” *American Philosophical Quarterly* 56:125–134.
- Fine, Kit. 2012. “Guide to Ground.” In Fabrice Correia and Benjamin Schnieder (eds.), *Metaphysical Grounding*, 37–80. Cambridge University Press.
- Korsgaard, Christine M. 1989. “Kant’s Analysis of Obligation.” *The Monist* 72:311–340.
- Markovits, Julia. 2010. “Acting for the Right Reasons.” *Philosophical Review* 119:201–242.
- Nelson, Mark. 2015. “What the Utilitarian Cannot Think.” *Ethical Theory and Moral Practice* 18:717–729.
- Olsen, Kristian. 2014. “Ross and the Particularism/Generalism Divide.” *Canadian Journal of Philosophy* 44:56–75.
- Parfit, Derek. 1984. *Reasons and Persons*. Oxford: Clarendon Press.
- Pettit, Philip and Smith, Michael. 2000. “Global Consequentialism.” In Brad Hooker, Elinor Mason, and Dale Miller (eds.), *Morality, Rules, and Consequences*. Edinburgh: Edinburgh University Press.
- Railton, Peter. 1984. “Alienation, consequentialism, and the demands of morality.” *Philosophy and Public Affairs* 13:134–171.

-
- Ross, David. 1930/2007. *The Right and the Good*. Oxford: Clarendon.
- Sharadin, Nathaniel. 2019. “Consequentialism and Moral Worth.” *Utilitas* 31:117–136.
- Smith, Michael. 1994. *The Moral Problem*. Blackwell.
- . 1996. “The Argument for Internalism: Reply to Miller.” *Analysis* 56:175–184.
- Stocker, Michael. 1976. “The schizophrenia of modern ethical theories.” *Journal of Philosophy* 73:453–466.
- Stratton-Lake, Philip. 2011. “Recalcitrant Pluralism.” *Ratio* 24:364–383.
- Svavarsdottir, Sigrun. 1999. “Moral Cognitivism and Motivation.” *Philosophical Review* 108:161–219.
- Williams, Bernard. 1981. “Persons, Character and Morality.” In *Moral Luck: Philosophical Papers, 1973-1980*. Cambridge University Press.
- Wood, Allen. 2017. *Formulas of the Moral Law*. Cambridge University Press.