
Note: This is the version of the paper submitted to *The Journal of Philosophy*, where it's now forthcoming. The published version is set to have some significant differences from this draft.

Against SCHMOUGHT

Abstract: Matti Eklund has argued that a new problem in metanormative theory arises when we consider the possibility of “normative counterparts”—normative concepts with the same normative roles as OUGHT and RIGHT (for instance), but with different extensions. I distinguish two versions of the problem, and propose a response: when we attend to the commitments involved in the possession and application of some normative concepts, we find that the possibility of normative counterparts is rationally ruled out.

Word Count: 7998 (with abstract, footnotes, and references)

1. Introduction

1.1. “*Evil Genius*”

The latest epic battle of good and evil comes, thank goodness, to a happy end. The evil plans of Evil Genius have been once again uncovered, his henchmen subdued, his diabolical Chaos Machine blown to smithereens. Having averted doomsday, our valiant Hero approaches her defeated nemesis, searching for adequate words of condemnation.

“You ought to stop trying to destroy the world,” Hero tells him sternly, after a pause. But Evil Genius doesn’t flinch.

Many thanks to Ray Buchanan, Jonathan Dancy, Andrew del Rio, Sinan Dogramaci, Julia Driver, Brigitte Gill, Matthew Matherne, Kent Mussell, Karim Nader, John Pittard, Brian Pollex, Anne Quaranto, Henry Schiller, Miriam Schoenfield, and David Sosa for their generous attention to the ideas in this paper.

“That’s true,” he grants, holding out his wrists for the handcuffs. “Still, you’ve got to admit, I *schmought* to keep on trying to destroy it, where ‘schmought’ is the action-guiding concept we in the criminal mastermind community use. In fact, while we’re playing the blame game, you certainly schmought not to have stopped me. That was, as we would say, very ‘schmrong’ of you.”

Hero never comes to regret saving the earth and its inhabitants from vaporization; but, long after she has handed Evil Genius over to the authorities, she does catch herself ruminating on his defiant rejoinder, and wondering: could he have been *correct* about what he schmought to do? And if he was, is there any way in which she herself, thinking in terms of what he *ought* to do, was *more* correct, more on the side of reality?

1.2. “Christmas Bonus”

Take another case: Hero receives a Christmas bonus. The thing is, though, she doesn’t really need the cash to get by, and her sensitive conscience is troubled: ought she to give away windfalls like this to people who are really struggling? Then a disturbing thought occurs to her. Even if, indeed, she *ought* to give the money away, might there not be some other, very similar standard applying to actions—call it the standard of what she “ought-lite” to do—that’s just a little more easy-going, so that it doesn’t include sacrifices like the donation of Christmas bonuses? And if so, why would she be so hell-bent on doing what she ought, instead of what she ought-lite? Don’t get her wrong: before this, she’d always thought it was very important to do what one ought. It’s just, now that there’s another standard for action that’s so similar—you ought-lite not kill people needlessly, for instance, or dump contaminants in rivers, and you even ought-lite to make charitable donations if you’re richer than Hero is—she doesn’t see much to

choose between. Still, it's no small thing to find your normative commitments in flux. Instead of feeling liberated to buy the new TV she wants, Hero feels confused.

1.3. Normative Counterparts and Normative Vertigo

The similarity between these cases lies in the strange turn each one takes when interrupted by what appears to be a certain kind of alternative concept: namely, one that shares a “normative role” with one of our normative concepts—it's used in the same way to guide or evaluate attitudes and actions—while having a different extension. Matti Eklund (2017) labels such role-identical, non-coextensive concepts “normative counterparts.” The prospect of these concepts, he argues, presents us with a new metanormative problem. As I'll interpret that problem, they seem to pressure us toward what we might think of as “normative vertigo”—a disorienting detachment from normative and evaluative facts, which suddenly look weirdly insubstantial.¹ And sure enough: in both our example cases, an encounter with putative normative counterparts leaves Hero feeling that the facts about what one ought to do have been somehow deflated. (In the next section I'll explore the different ways this happens in each case.)

This paper offers a solution to the problem of normative counterparts: I argue that, for some concepts, they don't exist—or, more cautiously, that we're rationally required to deny their existence. This is because to allow for such a possibility would be to incur rationally incompatible attitudinal commitments. If that's right, I'll argue below, then Eklund's problem disintegrates under scrutiny.

Here's what's coming. In section 2 I set out two distinct challenges that normative counterparts seem to raise. In section 3, I lay the groundwork for my solution by sketching a

¹ A distinction is often made between normative concepts like OUGHT or RIGHT and evaluative ones like GOOD or BEST. I'll ignore this distinction, however. The problem I'm interested in applies to both sorts of concepts (cf. Eklund 2017: 1, n. 1).

certain normative concept, calling it OUGHT. (I'll use small caps to refer to concepts.) Having introduced it, I show in section 4 why we have to deny that this concept could have any normative counterparts. Finally, section 5 explains the distinctiveness of my approach, and tries to dispel worries about it. We find ourselves in strange dialectical territory, when even our normative concepts have been called into question, and one might doubt that it can do any good to defend them. But I'm optimistic.

2. Challenges from Normative Counterparts

2.1. Setup: Plenitude and Parity

We can set up the problem of normative counterparts in two stages: call them *plenitude* and *parity*. First, plenitude. Once we allow for normative counterparts at all, there seems to be no principled way of keeping them from proliferating. There is, after all, a bottomless supply of extensions to build them out of—every combination of possible actions you can imagine. People may never have adopted normative concepts that possess those deviant extensions; but it's hard to see why such concepts *couldn't* be used. So it looks like we'll have to think there's a very large number of normative counterparts on the shelves.

That would mean that, wherever our normative concepts are used, we can be sure that there are counterparts pointing in every direction. You hold that the electoral college is an unjust institution. You tell your friend she should go vegan. You think *The Godfather* is the best movie of all time. And suppose all that's right. Still, if you have to allow for indefinitely many normative counterparts of UNJUST, SHOULD, and BEST, then other judgments can also be made. The electoral college is after all not *schmunjust*; your friend *shouldn't** go vegan; and the *best-prime* movie is actually *Casablanca*. And here's the problematic parity between normative

counterparts: these latter judgments are just as true as yours. The person who makes them is getting things right.

Importantly, this parity between normative counterparts can't be dissolved by appeal to higher-order normative truths. It's natural, for example, when first confronting the concept BEST-PRIME, to suspect that it is in some way a *bad* concept; that we *should* use trusty old BEST instead. And that may be true. Saying so, however, seems only to bring other normative concepts into play; and, once we've allowed for plenitudinous normative counterparts, we're vulnerable to a simple iteration of the problem. A lover of *Casablanca* might cheerfully agree that BEST-PRIME is bad; but suppose he also insists that our BEST is BAD-PRIME; we SHOULDN'T-PRIME concern ourselves with it. We've seen no reason, so far, that this couldn't be true. And the same thing might happen again one level higher, if we compare BAD and BAD-PRIME; and so on. The parity between concepts might go "all the way up." We can begin to explain the problem of normative vertigo by saying that such a robust parity between normative counterparts would be an uncomfortable one, what McPherson (forthcoming: 3) calls an "alarming symmetry."

2.2. *The Challenge to Ardent Realism*

Why uncomfortable? Why alarming? It may be useful to distinguish two separate challenges that normative counterparts can pose. The first—and the one Eklund gives the most attention to—is a challenge to a form of robust metanormative realism, what Eklund calls "ardent realism." *Ardent* realists are hardcore realists. They resist the minimalistic deflation of normative truths that quasi-realists go in for; they see questions about what to do as having some really, stoutly objective right answer; they hold that in some way "reality itself favors certain ways of valuing and acting" (Eklund 2017: 1). Before Eklund, it might have seemed enough to satisfy an ardent realist if she could show that there are bona fide facts (sufficiently explanatory

ones, maybe²) about what to do—for it to be mind-independently true, say, that Evil Genius ought to lead a reformed life of charitable work and gardening. But Hero’s experience in that first example case suggests otherwise. Normative counterparts can make us worry that such facts aren’t enough to vindicate the ardent realist’s belief in some kind of supreme normative standard.

The worry isn’t exactly that the existence of normative counterparts like SCHMOUGHT would show that normative judgments aren’t *objective*. No: in the “Evil Genius” case described it’s objectively and absolutely true that Hero’s nemesis ought to stop plotting mass murder. The problem is rather that this objectivity seems deflated—not what ardent realism was after. In fact, there are now too *many* objective truths: for, if we accept Evil Genius’s normative counterpart, it’s just as objectively true that he schmought to continue in his wicked ways. If that’s so, then it doesn’t seem like Hero’s practical stance is any more aligned with reality’s preferences (so to speak) than Evil Genius’s: as far as supervillainy goes, it looks like reality has declined to take a position of its own. Thus, without actually contradicting any realist normative claims, normative counterparts push us toward a less ardent realm: a realism unsettled by normative vertigo.

2.3. *The Challenge to Normative Resolve*

Even if the fate of ardent realism doesn’t interest you, the alarming symmetry of normative counterparts presents a further challenge, one that Eklund doesn’t emphasize: call it the challenge to normative resolve. This is what the “Christmas Bonus” case is meant to illustrate. There, Hero isn’t directly concerned about whether *reality* is, as it were, on the side of OUGHT rather than OUGHT-LITE. What she’s wondering about is whether to be on the side of OUGHT, *herself*, now that a normative counterpart is on the table. More generally, with such alternatives available, one’s commitments to respecting any normative property are in question.

² See Fine (2001: secs. 6–7) and Dreier (2004: sec. 6), who propose to distinguish realists from anti-realists and quasi-realists by the role they give to normative properties in certain grounding explanations.

How so? What's the worry this time? Well, it isn't exactly a worry that our normative concepts aren't significant, or that we should act in accordance with different ones (see Eklund 2017: vii). For the worry wouldn't be put to rest if we established that our normative concepts actually *are* significant, and are the ones that should guide our actions: that would only trigger a higher-order worry about SIGNIFICANT and SHOULD, next to their own normative counterparts. But if the objectivity of normative facts seemed deflated, before, in the challenge to ardent realism, now even the *significance* of these facts seems insubstantial and beside the point. If there's no more to say for OUGHT than that it's significant, just as it can be said for OUGHT-LITE that it's significant-lite, then a resolute commitment to either concept over the other might seem irrational. Or, even if strictly speaking such a commitment doesn't fall into the extension of IRRATIONAL, it nevertheless feels unattractive in some quasi-irrational way. That's the feeling of normative vertigo.

2.4. Comparison: The Normative Question

In this second form, especially, Eklund's problem can helpfully be compared to another criticism of normative realism (especially non-naturalism): the charge that it's unable to answer what Christine Korsgaard (1996: ch. 1) called "the Normative Question." James Dreier's (2015) and Shamik Dasgupta's (2017) and (2018) press challenges of this variety. They begin by granting to the non-naturalist that there's some *sui generis* property of the kind she makes goodness or rightness or being-a-reason out to be: call it property *P* (see Dasgupta 2017: 300). But they then pose an explanatory challenge: why should we accord special respect to *P*, over other properties, by calling it (say) *goodness*? Or, why would beliefs about *P*, instead of about all the other properties in the world, rationally require us to be motivated to act in certain ways? The non-naturalist realist has fixed on this *sui generis* feature of things, and she tells us it should

occupy a central place in deliberation and evaluation—but can she explain what it is about this feature that makes it so special? The challenger thinks not: any non-naturalist answer will fail to satisfy.

This problem, like Eklund’s, involves conceding something to the realist, and then raising doubts in one way or another about the importance of what has been conceded. A difference between the problems, though, should also be observed. Briefly put, Eklund concedes one point more than objectors like Dreier and Dasgupta do. They were ready to go as far as supposing that there was *something* of the kind the realist posited: a non-naturalist property, a “*sui generis* whatnot” (Dasgupta 2017: 301, paraphrasing David Lewis). But if it was also going to be maintained that this property *ought to be promoted*, or *rationally required* agents to respect it, or just *was goodness*, then they wanted to see some explanation of this. That last point, it seemed to them, was more than the non-naturalist could be allowed. But even after that point is conceded, it’s still possible to feel the force of Eklund’s problem, which is not an explanatory burden but the looming threat of normative counterparts. That threat might still make you queasy, even though, as we’ve seen, it’s not easily to formulate the problem precisely. I’ve spoken of it as a kind of affliction, even, as “normative vertigo.” There do seem to be philosophical grounds for the affliction, however, as this section has tried to bring out. It would be nice to have something to say to address it. Well, hang in there. I’m here to help.

3. Normative Roles and the Concept OUGHT

3.1. The Plausibility of Normative Counterparts

Take a step back. Why would the existence of normative counterparts seem plausible in the first place? Why should we expect that linguistic communities would be able to invent their

own versions of normative concepts, with different extensions than ours have? One reason is that some metanormative views seem to guarantee this possibility (see Eklund 2017: 20–22). Apart from that, though, I suppose there just don't seem to be many hard-and-fast constraints on these features of concepts: both extension and normative role look like simple matters of convention.³ If enough people think of drinking fountains as “bubblers,” then drinking fountains will indeed belong to the extension of BUBBLER. Or, if they persistently use a concept to, say, blame wrongdoers, or bestow honor, their concept will indeed *have* that function. Just mix and match role-conventions with extension-conventions, then (it seems) and you'll have a straightforward method for constructing normative counterparts. That, however, is the picture I'll be disputing. When we pay closer attention, we'll see that it leads to unacceptable consequences.

It won't try demonstrating *all* our normative concepts are counterpart-free. I'll be content if I can establish the point for one example case: the case of a concept I'll call OUGHT. If my strategy works here, it might work more broadly; and regardless, just a single concept without normative counterparts would be enough to break up the conceptual parity that Eklund's problem presented. Indeed, even a *possible* counterpart-free concept would be enough to answer his challenge (see Eklund 2017: 23). Therefore, I feel at liberty to make my consideration of OUGHT a little conjectural. I won't deliberately make things up. I'm proposing what seems to me a plausible account of how we use a certain real-life concept. But if you're skeptical about that, consider a more modest position. See if you think it's at least possible for a concept to work this way.

3.2. *The OUGHT-User's Commitments*

³ Thanks to Henry Schiller for helpful conversation in this area.

I'll characterize the concept OUGHT in terms of the conceptual role it plays. And this conceptual role, I'll suppose, can in turn be understood as consisting in the characteristic pattern of commitments that competent OUGHT-users undertake—here I'm taking cues from Ralph Wedgwood (2007), although the details of our proposals diverge. One of those commitments can be represented as follows, where S is an agent and φ is an action:

(1) The OUGHT-user commits to approving of S 's φ -ing if and only if S ought to φ .

(The OUGHT-user's commitment should be read as taking a wide scope over the biconditional.)

The idea here is that whether or not to approve of an agent's action won't be an open question for a competent OUGHT-user, once she judges whether the action is what S ought to do. If that attitudinal question *is* still open for her after that, it only shows she hasn't quite got the hang of the concept.

A person needs to have the commitment described in (1) in order to count as even possessing this concept. Actually deploying the concept will impose new, related commitments—most simply, perhaps, when one judges that S ought or ought not to φ . These judgments, I posit, involve undertaking commitments according to the following patterns:

(2) If you affirm that S ought to φ , then you're committed to approving of S 's φ -ing.

(3) If you deny that S ought to φ , then you're committed to refraining from approving of S 's φ -ing.

Think of these commitments, again, as a matter of settled questions. Once someone has decided whether S ought to φ , she won't still be in doubt about whether to approve of that action; not if she's manifesting competence. That question has already been settled—perhaps as the competent user of MAMMAL, when she applies that concept to an object, has already settled the question of whether it's an animal.

Now, this understanding of OUGHT, as a matter of committing to some non-doxastic attitude, might seem to have a non-cognitivist sheen to it—dubious comfort to the ardent realists Eklund was goading. However, self-described “robust realists” like Wedgwood and David Enoch (2011: 177–84) have adopted it, on the idea that they can still give normative terms suitably realist semantic values (see Wedgwood 2007: 106). Since Eklund, too (2017: 40–43; forthcoming: 8–9), treats the approach as ardent realism’s best answer to his problem, I won’t worry about its realist credentials. And I won’t linger longer, here, to fill out this rough conceptual sketch, which I hope is good enough for going on with. I’ll elaborate further only in stating two assumptions I make about the kind of question-settling commitments I’ve invoked. First, I’ll assume that the OUGHT-user undertakes these commitments whenever she competently applies the concept to actions. They aren’t like conversational implicatures she could evade by prefacing her predications with the phrase “strictly speaking,” or by making her judgments only privately. Second, I’ll assume that it’s irrational to commit both to approving of an action and to refraining from approving of it. I don’t mean for this to be a mere stipulation, though: I’m hoping you can already see the ugliness of settling an attitudinal question in both ways at once—not very different from the ugliness of believing or intending contraries.

3.3. Accepting the Concept

I’m officially agnostic about whether the concept I’ve just delineated is one we actually use, which affords me protection from certain semantic objections. But that doesn’t mean it’s beyond criticism in principle. Even if all I’m doing is proposing rules for a brand-new concept, there’s still reason to reject it if those rules would get us into trouble. That’s one lesson of A. N. Prior’s mischievous logical operator “tonk” (Prior 1960), the introduction and elimination rules of which would permit its user to reach any conclusion from any premise. Racial slurs would be

similarly defective concepts, on the inferentialist analysis Michael Dummett gives them: their rules of use license an inference from a person's racial identity to contemptuous beliefs or other pejorative attitudes about that person (Dummett 1973: 454).⁴ One way of explaining what's wrong with such concepts—following Belnap (1962: 132) and Dummett (1973: 397, 454)—is to say that TONK and racial slurs fail to be “conservative.” If we added them to our conceptual inventory, we would be adopting rules of inference that convey us to previously unlicensed conclusions: to grotesque and vicious errors, in fact. So, we need to be careful what concepts we pick up.

However, my OUGHT isn't like those concepts. The bare-bones patterns of commitment I've enumerated for it look simple and conservative: they're not going to force anyone into new, unwelcome approvals and non-approvals they hadn't counted on. If, in responding to Eklund's problem, I want to use a concept with such rules, I see no reason I can't.

4. Ruling Out Normative Counterparts

4.1. *The Strategy Here*

Now for the main event: I show why it's irrational, for any user of the OUGHT I've been describing, to accept the existence of normative counterparts to that concept. We're rationally required to reject such counterparts, because allowing for them would entangle us in an inconsistent set of commitments. This part of the argument comes in two phases. In Phase 1 I'll contend that it would be irrational to judge, of any *particular* action, that it falls within the extension of OUGHT but outside the extension of a normative counterpart. That gets us partway there. In Phase 2 I argue for a further conclusion: we are required to make a negative existential

⁴ To be clear, I have nothing staked on Dummett's analysis being correct, though it does matter that it could describe some *possible* concept.

judgment that there just *aren't* any normative counterparts of OUGHT. And if that's what we should think, then we shouldn't feel threatened by SCHMOUGHT, or OUGHT-LITE, or OUGHT*. We shouldn't so much as believe in them.

4.2. Phase 1: No Particular Divergences

To see why these judgments are rationally required, look at what it would mean for someone to fail to make them. Let's take Hero again as our protagonist, and suppose she rashly ignores the advice I'm giving, for some particular action: your adopting a puppy, say. Suppose Hero sincerely affirms that this puppy-adoption falls into the extension of OUGHT but not that of some arbitrary normative counterpart OUGHT*.⁵ If she does this, it amounts to affirming that you *ought* to adopt the puppy, while denying that you *ought** to do so. But in doing both those things, I'll now argue, Hero is either manifesting semantic incompetence or making a rational mistake.

Recall that normative counterparts, while having different extensions, must have the same normative role. (That's part of what it *is* to be a normative counterpart.) This means that our candidate normative counterpart, OUGHT*, has to function in just the ways I described in section 3, when talking about how OUGHT works. If a sincere and competent OUGHT*-user affirms that one *ought** to φ , she is committed to approving of that φ -ing; and if she denies that one *ought** to do something, she is committed to refraining from approving of the action.

Poor Hero, it follows, has undertaken rationally incompatible commitments. First, by adding your puppy-adoption to the extension of OUGHT, she has affirmed that you *ought* to adopt the puppy. From the normative role of OUGHT, we know she has thus committed herself to approving of the puppy-adoption—see commitment pattern (2). Second, though, in excluding the

⁵ I treat OUGHT as a monadic predicate with an extension composed of actions (all the actions whose subjects ought to perform them), but this is only for simplicity. To respect the grammar of the term in English, we'd want to find another logical form. See Schroeder (2011) for a roundup of some options, and an argument for thinking of it as a relation between agents and actions.

action from the extension of OUGHT*, she has denied that you *ought** to adopt the puppy. Thus, in accordance with OUGHT*'s normative role—its version of commitment pattern (3)—she has committed herself to *refraining* from approving of your puppy-adoption. That's the case, at least, if she's using the concept competently. Now put these two commitments together: Hero has committed herself both to approving of your adopting the puppy, and to refraining from approving of it. She has settled the question of approving of it both positively and negatively. But that's not okay. As I said above, I'm assuming this combination of commitments is irrational.

As one way drawing out the irrationality, imagine running the following experiment. You sit a volunteer down at a desk with a printed list of possible actions and a pen, and give him these two directives: "Put a check mark next to just those actions which you approve of. Also, *circle* just those actions which you approve of." You leave him to it. Later, when you come back to check his work, you find that your action of adopting the puppy has a check mark next to it, but isn't circled. What does this show? Well, your volunteer might just have made a mistake, manifesting less than complete competence in following the instructions. But if it's not that—if he really was deliberately, sincerely expressing both approval and non-approval, when he got to that item of the list—then something really strange has happened. He has suffered an unusually severe rational lapse.⁶

A rational lapse of this sort would be involved in granting that some action falls in the extension of OUGHT but not OUGHT* (or vice versa, as it's easy to see). Contemplating any case at all, you're rationally bound to assign it either to both extensions, or to neither; and this offers protection against normative vertigo in particular cases of confrontation, like the sort we see in

⁶ Miriam Schoenfield helpfully pushed me to consider the *degree* of irrationality involved in these patterns of commitment.

“Evil Genius.” If Hero is right to be sure that her nemesis ought not vaporize the world, then she’s also right to be sure that he schmought not vaporize the world.

4.3. Phase 2: No Normative Counterparts

Hero may have other worries, though. Maybe she hasn’t found any cases where what someone ought to do comes apart from what they ought* to do, but she might wonder all the same if those cases are out there. Can we help her shake that nagging thought? Can we show it’s rational for her not just to *avoid* assigning *different* extensions to OUGHT and OUGHT*, but to believe that such role-identical concepts have the *same* extension? I think we can.⁷

Attend to the biconditional commitment involved in the possession of the concept OUGHT: a competent OUGHT-user, says (1), is committed to approving of S’s φ -ing if and only if S ought to φ . Since OUGHT* is meant to be a normative counterpart of OUGHT, we have to take it that it carries the same biconditional commitment: the competent OUGHT*-user must commit to approving of S’s φ -ing if and only if S ought* to φ . We also suppose that Hero is competent with both concepts: she understands the normative roles they play, at least implicitly. We can then ask: when, if at all, will it be rational for Hero to undertake *both* the biconditional commitment of OUGHT, and that of OUGHT*? And the answer must be: only so long as she regards them as coextensive. Thus, she can’t think they’re normative counterparts after all: by definition, normative counterparts have different extensions.

Now, why must this be Hero’s answer? Well, suppose otherwise. Suppose there’s even one action—let it be that puppy-adoption of yours, again—such that Hero is agnostic about whether it might be sorted into OUGHT’s extension but not OUGHT*’s. As a competent OUGHT-user, Hero is committed to approving of your adopting the puppy just in case you ought to adopt

⁷ In this section I’m indebted to Sinan Dogramaci, who helped me especially in thinking through Phase 2’s argument.

it. As a competent OUGHT*-user, Hero is also committed to approving of adopting the puppy just in case you *ought** to adopt it. Now, notice that these two conditional commitments will turn out to be incompatible, unless the puppy-adoption's status with respect to OUGHT is the same as its status with respect to OUGHT*. But then, if Hero is agnostic about whether these concepts are coextensive, then what she's doing is irrational. Its irrationality is something like the irrationality of judging that something is a puma just in case it's a mountain lion, and judging that it's a puma just in case it's a cougar, while being agnostic about whether PUMA and COUGAR are coextensive. Such a set of conditional commitments would involve, again, a rational failing of uncommon magnitude.

Another illustrative experiment you can run: Set two lists of actions on the desk in front of your volunteer. Gesture to the first, and ask him, "Do you approve of just the actions that appear on this list?" If he says yes, gesture to the second list. Ask: "Do you approve of just the actions that appear on *this* list?" If he says yes again—and if he understands what he's doing—then either he believes that the lists contain all the same actions, or else he's haphazardly settling attitudinal questions in a way that may well be incompatible. But that, I take it, would be irrational.

4.4. *Lessons Learned*

OUGHT* stood for any posited normative counterpart of OUGHT: the conclusions we reached with it have broad implications. It doesn't make sense, we've found, to allow that *any* concept sharing OUGHT's normative role could differ with it in extension: that would make your set of conceptual commitments incoherent. And so we've arrived at the main point of this paper: for the sincere OUGHT-user, toleration for normative counterparts is off limits. To avoid attitudinal inconsistency, she must instead reject them whenever they're proposed—and that

requirement is good news, I take it, when we come to Eklund's problem. It means that we should never get to the point of feeling the quasi-rational pressure toward normative vertigo that we invite if we take on board an endless supply of normative-conceptual options, all structurally on par with each other. The normative-conceptual field has been winnowed, the grounds of normative vertigo removed.

One upshot of this argument is that the naïve picture of concepts' normative roles and extensions we started with, in section 3.1, was flawed. Let enough people use a concept for a certain purpose, we might have thought, and let them do it with a certain group of referents in mind, and they'll eventually succeed in combining that role and that extension. We'd have to grant this, it might have seemed, even if we ourselves don't bother using their concept. But that's going too fast. An affirmation that some object is in the extension of a concept isn't just neutral metasemantic record-keeping. It also amounts to a *use* of the concept, or at least an affirmation of its use. Therefore, we need to be circumspect about the doxastic or attitudinal inferences that the concept's use would commit us to: as with TONK and (according to Dummett) racial slurs, so also with SCHMOUGHT.

In a funny way, then, Eklund's problem and its solution go hand in hand. For the problem wasn't just that we seemed bound to acknowledge lots and lots of concepts, each with a different extension. That's painless enough, on its own. The thing was, all the different concepts Eklund was conjuring up were competing for the very place in deliberation or evaluation that our more familiar normative concepts held. That is, it was the sameness of normatives role that made normative counterparts feel so disturbing. But, I've now argued, it's this very sameness of the normative roles, with their patterns of commitment, that protects the alert OUGHT-user from vertigo. If she's being rational, she must judge that her concept has what Eklund calls

“referential normativity”: its normative role fully determines the extension of any concept associated with it (2017: 10).

5. Defending My Metasemantic Approach

5.1. My Contribution

In Eklund’s estimation as in mine, a defense of referential normativity—a rejection of normative counterparts—is the most promising way of responding his problem. He even points to Wedgwood’s conceptual role semantics, which I’ve been riffing on, as a way of securing that thesis (Eklund 2017: 40–43). So, he wouldn’t find my form of metasemantic response entirely unexpected. In this paper, though, I’ve done more than just to elaborate Eklund’s own tentative suggestion: I’ve offered a more economical solution. It’s true that a Wedgwoodian account of normative concepts would guarantee that concepts with the same roles get matched with the same extension; but that account also depends on some substantial positions I haven’t needed to endorse, myself. For instance, it requires there to be constitutive goals of theoretical and practical reasoning (see Wedgwood 2007: 51–52; 100–102), and it rests on a general theory of reference determination based on preserving the validity and completeness of a concept’s rules of use (Wedgwood 2007: 99; cf. Peacocke 1992: 19).⁸ My premises and conclusion are compatible with these assumptions, but, for my part, I’ve relied only on the ways in which a possible normative concept could function. As a result, my answer to Eklund’s problem has remained pleasantly lightweight.

My approach has also offered an insight into the status of referential normativity. It shows that this metasemantic thesis isn’t just a desperate tactic for ardent realists on the ropes.

⁸ Thanks to Ray Buchanan for valuable discussion here.

No, if I'm right, it's required of *everyone*—at least, everyone willing to accept the concept I sketched in section 3. Just by being an OUGHT-user, I said, one already undertakes a certain set of attitudinal commitments; and section 4 argued that these commitments rationally require one to deny the existence of normative counterparts. In fact, then, even if you yourself never felt particularly disturbed by those alternative concepts—some philosophical constitutions bear them more easily than others—you, too, must reject them. Complacency here is irrationality.

I take that to be a significant finding. The rational requirement that *everyone* affirm referential normativity, for some concepts, is a noteworthy feature of the metanormative landscape, and would be worth pointing out even if it didn't also help defuse the concerns Eklund raised. But, as I've argued, it does that too.

5.2. *The Embarrassment of Riches Objection*

There is this further advantage, too, in my approach to Eklund's problem: it extends straightforwardly to address his chief worry about solutions from conceptual roles, the "very serious objection" he calls the "embarrassment of riches" (Eklund 2017: 54). Even if concepts with the *same* normative roles must have the same extension, Eklund asks, couldn't there be alternative normative concepts with just *slightly different* normative roles? (We can call these concepts "normative relatives.") For instance, I've understood OUGHT's normative role in terms of commitments governing one's approval. Could there be a community that instead uses the concept OUGHT-ISH, which involves commitments governing the similar attitude "approval-ish"? Maybe this is like approval but with slightly shifted boundaries, a different precisification of a vague phenomenon. (See Eklund 2017: 56.)

If normative relatives are possible, nothing will force them to match our concepts in extension. One community might truly affirm that S "ought-ish" to φ , while we truly deny that S

ought to φ . And if perhaps these judgments don't conflict as obviously as judgments involving exact normative counterparts, they still seem ill at ease. Accepting them both would commit us to approve-ish-ing of something we don't approve of. And if approval-ish is enough *like* approval, that looks bad: OUGHT and OUGHT-ISH would appear to be in "normative competition" very similar to the kind we saw between normative counterparts (Eklund 2017: 55). This conflict might also be repeated at higher orders, just as we imagined with normative counterparts: they would then present us with "the same sort of parity problem" we saw before (Eklund 2017: 54), and threaten to induce a closely related form of normative vertigo.

I can't show that normative relatives of OUGHT have to be considered coextensive with it. We wouldn't want to conclude that, anyway, without knowing more about the particular role-variations they involve. However, the considerations I brought out before do give us reason for optimism about this mutation of Eklund's problem: here, too, we can dare to hope, the problem and the solution go hand in hand. For, first, we know it isn't the sheer plenitude of normative relatives that would feel disorienting, but normative competition between them. Challenges to ardent realism and normative resolve will only arise when two concepts would seem to commit us to rationally incompatible attitudes toward the same actions. But then the solution to the problem has the same form as it did for normative counterparts. In some cases, yes, OUGHT will prescribe one attitude toward some act, and accepting that a normative relative has a different extension would commit us to a rationally incompatible attitude about it—but in those very cases, rationality requires us not to assign such deviant extensions. And so the problem dissolves.

5.3. The Metaphysicians' Objection

Stephanie Leary and Tristram McPherson have argued that a metasemantic response to Eklund's problem, such as this defense of referential normativity, would miss the point, at least

as far as ardent realism goes. (Leave aside the question of normative resolve, for now.) They note that the guiding idea of ardent realism—or the featured metaphor, anyway—is that of reality itself favoring certain ways of valuing and acting. But then, to them it seems that this is simply a matter of there being some real normative properties out there in the world: a matter for metaphysics, not metasemantics, to arbitrate. If Evil Genius and Hero, say, are operating with different concepts, referring to different properties, then the question we need to ask is just “which properties are really normative” (Leary forthcoming: 7). That is, which of them has “that *sui generis prescriptivity* that objectively ‘calls out’ for certain responses in us”? (Leary forthcoming: 11; original italics). And if one of the properties is *really*, intrinsically normative in this way, and the other isn’t, then that should be enough to show that there’s no problematic parity between the two, right? If that’s how things are, then “ardent realism appears directly secured, independently of referential normativity” (McPherson forthcoming: 13).

I’ll say this, for such an unruffled metaphysical response: it may be enough to answer the problem Dreier and Dasgupta were pressing (see section 2.4). They wanted, you’ll recall, an explanation of why some non-natural metaphysical property, some “*sui generis* whatnot,” would deserve the kind of respect we give to reasons or goodness. Leary and McPherson might reply: it deserves such respect because it is a *normative* whatnot; it’s a whatnot possessed of an intrinsic prescriptivity, objectively respect-worthy. Of course, that kind of flat-footed answer isn’t likely to win the skeptics over. Dreier will likely call it “unacceptably mysterian” (2015: 180); Dasgupta may judge it “a pernicious remnant of Scholastic metaphysics” (2017: 312). *If* you’ll grant the metaphysics, though, it seems the ardent realist can resist this objection.

But on my reading, remember, Eklund’s problem is different precisely in that it isn’t disposed of even by such generous suppositions. It’s no surprise, then, that I think the flatfooted

metaphysical response fails to resolve it. For, imagine we grant what the metaphysicians ask for: the property Hero picks out with OUGHT, and *not* the one Evil Genius picks out with SCHMOUGHT, is *really* normative. McPherson thinks this stipulation would rule out the hypothesis of alarming symmetry we were worried about: he thinks it's the hypothesis that there is "nothing that normatively distinguishes" the two properties (McPherson: forthcoming: 13). But that isn't the way I understand Eklund's problem (see section 2). Rather, the alarming symmetry would now manifest in the form of higher-order normative counterparts: we must imagine that while there is indeed a normative difference between the two properties, there is also a *schnormative* difference cutting the opposite way: suppose only SCHMOUGHT, not OUGHT, picks out a genuinely *schnormative* property. It alone has intrinsic *schmescriptivity*. If that's how things are, it would be strange for the ardent realist to feel vindicated. Eklund's readers are likely to feel light-headed again, for all that the metaphysicians have done to steady them.

This feeling of normative vertigo would certainly be a mistake, on my view: that's what I've been arguing all along. It did need to be argued for, though, if Eklund's problem was to be addressed, and to my mind McPherson and Leary haven't quite addressed it. The metasemantic style of response, on the other hand, seems to scratch the right itch.

5.4. *The SCHMATONALITY Objection*

Throughout this paper, I've talked about what rationality requires of the OUGHT-user. I've concluded, in the end, that she should deny the existence of normative counterparts. There is a reply to this that comes all too naturally, though, once you've read your Eklund. "Rationality may require the rejection of SCHMOUGHT," you might find yourself saying; "but what does *schmatonality* require? Or, what *should** we think of normative counterparts?" I've been

defending OUGHT, my flagship normative concept, from competitors; but I've used normative language in the process. Does that make me guilty of some sort of question-begging?

I think not. I have argued that it's irrational for an OUGHT-user to worry about SCHMOUGHT and other normative counterparts. Now, if someone asks her whether it would be *schmational* to worry about them, I recommend she do the same thing as for SCHMOUGHT itself: she should reject the concept she's being handed. At least, it seems up to her challenger to justify an expansion of her conceptual repertoire: in her current position, as a RATIONAL-user, SCHMATIONAL looks likely to be defective. And it wouldn't be fair, in response to this move, to accuse her of question-begging because of her presumption that RATIONAL matters. That's what she came into the conversation thinking, after all; and it's to people with such presumptions that Eklund addressed his challenges. If those challenges have been based on concepts she's (rightly) suspicious of, it isn't *her* dialectical blunder that vitiates the discussion.

Maybe there's a different worry behind the objection, though.⁹ Even if my argument isn't exactly begging the question, you might wonder who it's *for*. Can anyone benefit from it? If some reader of Eklund wasn't bothered by normative counterparts, then this argument will be uninteresting to her; and if instead she was knocked off balance by them, won't she be too hesitant about the significance of rationality to follow the line of thought I've presented here?

I answer: that depends on the reader, but I remain hopeful. First, there's no cause for thinking that the only people worried by Eklund's problem have already ceased to privilege their familiar normative concepts. They might persist in normative resolve, even while recognizing a new quasi-rational tension in their belief and practice, or while losing some credence in ardent realism. The argument I've provided might help them resolve such tension or embrace that

⁹ In working through the dialectic, here, I was especially helped by conversations with Miriam Schoenfield and David Sosa.

metanormative view. Second, though, I do not abandon to their fate even readers of Eklund with more severe cases of normative vertigo: those who have gone so astray that they can no longer treat the facts about rational requirement as conclusive verdicts on attitudes or actions. Now, to be sure: if their deconversion has been so radical that *nothing* matters for them any longer, if they find themselves with no considered preference for even the most paradigmatically rational behavior over the most paradigmatically irrational, then they're beyond my power to help. May God have mercy on their souls. But they may not be like that. It may be that while the *label* "rationality" no longer moves them, they would still welcome some patterns of thought and behavior while finding others offensive: they're not ready to believe contradictions or shrug at global destruction. If there are Eklund-readers like this, then my argument might help to "scare them straight." A little experimentation with putative normative counterparts seemed innocent enough to them at first, in their naivete. But now I'm trying to show them the uglier side of their new lifestyle: if OUGHT was among the concepts they allowed counterparts for, and if they really understood what they were doing, then they were in fact committing a pretty grotesque rational error: at a first approximation, the error of both deciding to approve of something and deciding not to approve of it. Yikes. It isn't too late for them, though, even now. With open arms, rationality awaits their return.

6. Conclusion

Imagine that Eklund had mounted a complicated formal argument concluding that morality and rationality don't matter; but imagine this argument hinged on a hard-to-discern fallacy, or depended on premises that turned out to contradict each other. Had that been discovered, I suspect that worries about normative counterparts would have troubled no one;

those worries would have been dismissed as unmotivated. If Eklund's problem looked formidable, surely that's because it didn't seem to depend on such mistakes. It looked like worries about ardent realism and normative resolve could issue from clearheaded philosophical pondering. But in fact, I've been saying, the mirage of normative counterparts for OUGHT could appear only when we have suffered a rational lapse similar to the embrace of contradictory beliefs or intentions: it could arise just when we have incurred incompatible attitudinal commitments with respect to approving action. If she avoids that kind of error, then, Hero won't be daunted by a vast selection of normative concepts in every size and color. She won't fear that her practical deliberations have been trivially parochial. Add a *schm-* to any term you want, but she'll remain unruffled. To her, SCHMOUGHT is just her old friend OUGHT, charmingly misspelled.

References

- Belnap, Nuel (1962). “Tonk, Plonk, and Plink.” *Analysis* 22 (6): 130–34.
- Dasgupta, Shamik (2017). “Normative Non-Naturalism and the Problem of Authority.” *Proceedings of the Aristotelian Society* 117 (3): 297–319.
- Dasgupta, Shamik (2018). “Realism and the Absence of Value.” *Philosophical Review* 127 (3): 279–322.
- Dreier, James (2004). “Metaethics and the Problem of Creeping Minimalism.” *Philosophical Perspectives* 18 (1): 23–44.
- Dreier, James (2015). “Can Reasons Fundamentalism Answer the Normative Question?” In Gunnar Björnsson, Caj Strandberg, Ragnar Francén Olinder, John Eriksson & Fredrik Björklund (eds.), *Motivational Internalism*. Oxford: OUP.
- Dummett, Michael (1973). *Frege: Philosophy of Language*. London: Duckworth.
- Eklund, Matti (2017). *Choosing Normative Concepts*. Oxford: OUP.
- Eklund, Matti (forthcoming). “Reply to Critics.” *Inquiry*.
- Enoch, David (2011). *Taking Morality Seriously*. Oxford: OUP.
- Fine, Kit (2001). “The Question of Realism.” *Philosopher’s Imprint* 1 (2): 1–30.
- Korsgaard, Christine (1996). *The Sources of Normativity*. Ed. Onora O’Neill. Cambridge: Cambridge UP.
- Leary, Stephanie (forthcoming). “Choosing Normative Properties: A Reply to Eklund’s *Choosing Normative Concepts*.” *Inquiry*.
- McPherson, Tristram (forthcoming). “Ardent Realism without Referential Normativity.” *Inquiry*.
- Peacocke, Christopher (1992). *A Study of Concepts*. Boston: Cambridge: MIT Press.
- Prior, A. N. (1960). “The Runabout Inference-Ticket.” *Analysis* 21 (2): 38–39.

Schroeder, Mark (2011). "Ought, Agents, and Actions." *Philosophical Review* 120 (1): 1–41.

Wedgwood, Ralph (2007). *The Nature of Normativity*. Oxford: OUP.