White Rose Research Online

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

Deposited via The University of Leeds.

White Rose Research Online URL for this paper:
https://eprints.whiterose.ac.uk/id/eprint/135305/

Version: Accepted Version

Book Section:
Woods, J (2018) Footing the Cost (of Normative Subjectivism). In: Suikkanen, J and Kauppinen, A, (eds.) Methodology and Moral Philosophy. Routledge Studies in Ethics and Moral Theory, 1. Routledge, Abingdon, UK. ISBN: 9781138322097.

© 2019, Taylor and Francis. This is an author produced version of a book chapter published by Taylor & Francis in Methodology and Moral Philosophy, available online at https://www.routledge.com/9781138322097. Uploaded in accordance with the publisher's self-archiving policy.

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

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

Footing the Cost
(Of Normative Subjectivism)*

Forthcoming in *New Methods in Ethics*, eds. Jussi Suikkanen and Antti Kauppinen

Jack Woods

1. Introduction

According to some, there's rarely sufficient moral reason to interfere with reasonable adults' decisions about what to do since our lives are our own. According to others, there's often sufficient moral reason to do so since our choices impact on the welfare of others. And, of course, many of us vacillate between these viewpoints over the course of our lives. Many philosophers believe at least one of these viewpoints must be wrong—and wrong no matter who holds it and where it's assessed from. Moreover, many philosophers seem to think that believing that both could[1] be right—believing that some form of normative subjectivism is true—would undermine the point of normative judgment entirely. Philippa Foot reports the phenomenon thus:

> [subjectivism] may, as I said, appear dangerous and subversive of morality. We are apt to panic at the thought that we ourselves, or other people, might stop caring about the things we do care about, and we feel that the categorical imperative gives us some control over the situation. (Foot 1972)

*Thanks to Derek Baker, Catharine Diehl, Jessica Isserow, Alex King, Barry Maguire, Erum Naqvi, Ville Paukkonen, Paolo (Lil' P) Santorio, Karl Schafer, Lucas Thorpe, Teemu Toppinen, Ken Westphal, Gözde Yıldırım and audiences at Boğaziçi University and the University of Helsinki for useful comments. Thanks especially to Jimmy Lenman and Pekka Väyrynen who provided very useful comments in writing this final version. And, finally, thanks to my philosophical hero, Philippa Foot, for providing inspiration.

[1]Could how? Stay tuned.

and anyone who's taken a serious crack at developing subjectivism knows the reaction all too well.

The worry is not just that subjectivism is false. It's that subjectivism (or believing in it) might be *actively harmful* to the role that normative judgment plays in structuring our lives, evaluating others, planning our future, and the like. Given that this kind of complaint is rather woolly, I'll explore here a precise version: that somehow the truth or acceptance of subjectivism about normativity undermines the *functional role* of normative judgment. If subjectivism doesn't undermine the functional role of morality, as I'll claim, then the putative perniciousness that so many claim subjectivism brings in its wake is significantly diminished, if not entirely extinguished.

I claim that neither the truth of, nor belief in, subjectivism has deleterious impact on the functional role of normative judgments. There's no need to panic; at worst, the majority of the population simply believe falsely. I start by describing reasonable, yet unabashedly subjectivist positions (§2) and distinguishing two senses of normative universality (in §3). *Evaluative normative universality* holds that normative reasons are invariant under change in context when assessed from our actual normative standpoint. *Ontic normative universality* holds that normative reasons are invariant under change in context when assessed by the norms grounded in those very contexts.[2] I then argue that, for many aspects of the functional role of normativity, the putative costs of normative contingency are typically due to failures or perceived failures of evaluative normative universality. Yet subjectivism need not—in fact, for some subjects and some contexts, can not—deny evaluative normative universality.

Subjectivist views are ontically contingent by definition, so I need to show the costs of accepting ontic normative contingency aren't onerous. I address four ways in which ontic contigency might interfere with the functional role of normative judgment. In (§4-5), I argue that the ontic contingency of our reasons doesn't undermine their normative significance for us or for evaluating others. In (§6), I show that we can make sense of *why* we engage in normative communication—at least in many cases—given ontic contingency. (§7) explains how normatively evaluating hypothetical situations makes sense even given subjectivism. Finally,

[2]Compare Gibbard's (1990) distinction of *grandoise* objectivity from *moderate* objectivity for a similar distinction, albeit one in an expressivist context. See also (Lenman 2014) for discussion.

(§8) addresses whether certain intuitively false counternormative conditionals come out true on subjectivist views.

Admittedly, there are still features of subjectivism which conflict with intuitions about normativity. But the arguments and constructions of §4-8 alleviate the symptoms even if they don't cure the "disease". Much of the necessary palliative work I'll engage in is inspired by an insight of Foot (1972): our normative outlooks are robust in the sense that it would take a quite radical departure from ordinary humanity for our normative outlook to significantly change. Moreover, even if this isn't actually true, we believe it, and this belief stabilizes our normative practice. There's nothing magical here; just reasonable and plausible optimism about our similarity with one another.

I won't argue for the sociological claim directly, but there are good reasons to believe it: just look around and talk to folks for a while. Jimmy Lenman reports the thought nicely:

> The moral outlook we possess and express in our moral judgements has a generous measure of built in moral mind-independence. So, according to the outlook most of us share, the wrongness of killing wouldn't change if I stopped disapproving of it. But of course it's just our outlook. The moral truths it expresses did not shape its evolution. It is Man who has shaped the Way, the Way did not shape Man. (Lenman 2014: 243)

Given this thought, much superficial divergence in normative outlook is really *derivative* in the sense that some locally contingent non-normative fact explains the normative divergence. Typically, on inspection, our more-or-less shared normative outlook remains unshook.

A complementary point which plays a co-inspirational role comes from Barry Stroud's (1965) discussion of Wittgenstein's Wood Sellers—strange people who measure amounts of wood by the area instead of volume. Stroud argues that the way which we actually conceptualize Wittgenstein's example involves illicitly importing many of our own norms about how to measure. But evaluating such a case under this presumption does not show that their practices are irrational—it just shows that the irrationality of measuring wood that way while maintaining the remainder of our measuring norms.[3] We need to remember that when we consider

[3]This also echoes a useful discussion of similar issues by Sharon Street (Street 2009, lessons

bizarre scenarios we typically take our world with us. So it's not obvious that our initial reactions to such cases are probative; no one yet has fleshed out a counternormative case to the point where we could take it seriously as a challenge to serious subjectivist views.

Whether it's at all possible to conceptualize obeying such bizarre conventions from the *inside*, so to speak, is extremely fraught.[4] This point is somewhat known, but too infrequently recognized, as endless discussions of Parfit's (1984) future tuesday indifference example and frequent opposition to the rationality of Gibbard's (1990) totally coherent Caligula, etc. shows. Recognition of it diminishes the probative weight of putative counterexamples to normative subjectivism. As I'll argue below, this point is dramatized by the fact that we're quick to conflate evaluative contingency with ontic contingency when theorizing. This is extremely problematic when evaluating subjectivist views, given that we subjectivists need only recognize ontic contingency. Foot and Stroud's points, in tandem, do yeoman work in showing that we are far more capable of capturing the functional role and point of normative discourse than is typically thought.

My aim is to put together a number of existing points about subjectivism and functional role which, collectively, undermine the widespread sense that our normative judgments must be universal.[5] I agree that if we lacked evaluative normative universality, then we'd have a serious problem. But it's entirely unclear what the costs of ontic normative contingency are. This fact strengthens existing cases for subjectivist views—given, for example, by Foot (1972), Harman (1975), Wong (1984), Dreier (1990), Wiggins (1998), Sobel (2009), Velleman (2013), Finlay (2014), among others—by showing that costs of the view are relatively minor. There's no need to panic, even though subjectivism is true, even if we believe it.

4-5, 10-11) that we can't make changes in normative outlook without it spreading throughout the rest of our normative standpoint.

[4]Though her target is different, Foot's (1958) discussion of such examples makes this point vividly. As she claims, whenever we seem to understand how clasping one's hands three times could be basically good, there's typically a suppressed principle connecting this otherwise pointless activity to something obviously good like respect.

[5]Many techniques I use below are known in at least philosophical folklore; nevertheless, putting them together constitutes a novel treatment of this problem. Among those not otherwise cited below, the discussion of rigidified value judgments in Lewis (1989) is especially relevant and useful.

2. Normative Subjectivism

Subjectivist views, as I understand them, are characterized by an explanatory biconditional like:

> We have to reason to φ if and only if (and because) A

where A is some condition or state of affairs that is essentially dependent on *contingent* features of us, our community, our conventions, etc. Note that this biconditional is universally quantified and the 'because' here is meant as full, not partial, non-causal explanation. This characterization would need refinement to count as a definition of subjectivism—after all, the conjunction of a non-contingent normative principle and contingent facts about us shouldn't necessarily count as subjectivist—but I'll leave it at this slightly more intuitive level for now. The examples and discussion below should suffice to make clear what I have in mind.

Reasonable subjectivist views allow:

> It's correct *in some sense* to say, of a, that they have reason to φ if and only if and (because) A

where again A is a condition or state of affairs that is essentially dependent on contingent features of *me or my* moral community. This means that it's possible—though not required—that I, a member of one moral community, can correctly say of you, a member of another moral community, that you have reason to φ *even if* you can't say this of yourself.[6] I can correctly condemn the moral practices of another community even if their practice is not correctly condemnable from within. Not all subjectivist views accept this, but ones which don't have trouble providing any kind of universality. We'll henceforth put them aside as non-starters.

Our grounds for asserting that someone has a reason to φ need not be A (though they might be). Rather, the correctness of our assertion is explained by A, even if our grounds for asserting it aren't A. Compare a positivistic picture of the law where the explanation of why a particular fact is legal is grounded out in the decisions of the salient law-makers *and* our recognition of these law-makers as law

[6]For example, Lenman (1999) points out "However, [response-dependency] is consistent with my insistence that torturing human babies for fun is wrong applying not just to the actual world but to any world - for it may be part of the substantive content of my moral judgement that this wrongness is not conditional on my responses."

makers. This story about legal validity doesn't demand, and in fact shouldn't demand, that when we assert that something is legal, we take the grounds of the actual explanation as our justification for our assertion. That's an additional controversial commitment. Similarly, we might think that even though it's the general social goodness of relationships that make them good, we can't cite this fact to our partners to justify why we're nice to them.

Subjectivist views come in many varieties: non-quasi-realist expressivism,[7] Foot's (1972) instrumentalism, Street's (2010, 2012) Humean constructivism, my quasi-conventionalism (2016, forthcoming). To focus in, we'll take the following account of reasons as our working example—note that it's intended just as a working example—of a paradigmatically subjectivist view:

> CONVENTIONALISM: Someone has reason to φ in a situation γ just in case and because they're directed to φ in γ by a system of conventional norms which governs their potential action.[8]

This isn't a claim about the lexical *meaning* of reasons claims, but one about *what it is* for there to be a normative reason for me to do something. Given CONVENTIONALISM, suitable shifts in the norms we accept will generate a shift in the normative facts. However, even given CONVENTIONALISM, we needn't endorse claims like "If our norms permitted dog kicking, kicking dogs would be okay" (see §8). Rather, *qua* normative theorist, I endorse the *metanormative* claim that, if CONVENTIONALISM were correct, then if we were to have accepted norms permitting kicking dogs, we would have had (from the view of that context) reason to kick dogs.

This metanormative claim is crucially different from a similar sounding normative claim made from *within* our perspective. I put to the side here those, like Blackburn (1998) and Dworkin (1996), who suggest that we cannot make sense of the metanormative claim. Even if they're right about the particular metaethical views they favor, which I strongly doubt, we can make the necessary distinction for CONVENTIONALISM.

[7] Quasi-realist expressivism is more difficult to tackle. See Lenman (2014) for a version of quasi-realism which is explicitly not realism; others are less forthright.

[8] This way of articulating CONVENTIONALISM allows that they could also be directed to refrain from ϕ-ing. A full spelling out of this view needs to accommodate differential weights of reasons and directions in order to accommodate our complex normative conventions. We'll put this complication aside.

I won't chisel down the right-hand side of CONVENTIONALISM. We should add hedges to avoid worries having to do with false information, rash approvals, and the like, but presumably such chiseling can be done.[9] I talk in terms of reasons, not obligations, but the arguments work equally either way. I use reasons talk since it's increasingly pervasive and we can plausibly recapture obligation-talk in terms of having most reason. Anyway, though the contours of any actual subjectivist view will be tremendously complicated, this does not affect the general structural points I want to make.

For similar reasons, I'll put aside complicated relationships between moral normativity and our "thin" reasons to engage in moral behavior. For the sort of structural points I want to make here, these relationships don't matter much and my claims generalize across normative standards down to "thinner" normative notions. If anything, working with morality stacks the deck against me since it's exactly in the case of morality that our fears about deviant normative standards are greatest. For those wanting a plausible picture of how to merge two types of subjectivism to accommodate this distinction, I offer my (2016, forthcoming, manuscript) as examples.

Of course, to fully develop CONVENTIONALISM as more than an example, I'd need to give an account of what it is for a conventional system of norms to be in force for a particular agent, as I do in my (forthcoming). I won't repeat this story here since it's independent of my general claims. If a picture helps, it's close enough to think in terms of sophisticated positivism about the law (Hart 1961). I'll presume that conventional facts tell us both what reasons we have and what their strength is. This is a substantial assumption, but, again, it will do for here. Again, the points made below are general enough to hold regardless of one's normative topography. Finally, I have no truck here with *metanormative* subjectivism where what it is for there to be a reason at all, moral or otherwise, shifts from context to context, here being that we favor it, there being that God says that we must do such and so. That view dances far too close to literal incoherence, as Plato pointed out long ago. We turn now to distinguishing two senses in which we might think reasons are universal.

[9] I'll likewise not address the distinction between objective and subjective—in the sense of information-dependent—reasons below. A full account should, but I have no room for epicycles.

3. Two Notions of Normative Universalism

We may view normative reasons as holding for each other when *evaluated at our own perspective*. When you suggest you can ignore starving people and shop exclusively at Waitrose or Whole Foods,[10] I draw on my belief that we shouldn't harm innocents to persuade you to shop at Aldi's[11] and donate the difference to Oxfam. If you reject this reason, as it's no part of your (recognized) normative perspective, I still have no reason to revise my opinion that you're doing something you've significant reason not to. After all, I have strong evidence your normative beliefs are mistaken: my normative beliefs.[12]

Call this the *evaluative* sense of normative universality: that you (or anyone, but stick with 'you' for now) have reason to do something is evaluatively normatively universal when it holds, no matter what the non-normative facts may be, *when evaluated from our normative perspective*. For example, consider a normative perspective which only cares about utility maximization. For it, "we have reason to do the thing which is utility maximizing" is evaluatively normatively universal: no matter what conventions govern a context we evaluate, what matters is whether the actions performed there are utility maximizing. If their normative viewpoint tells them to do something which doesn't maximize utility, they're simply wrong (by our lights.)

Alternatively, we may view normative reason claims as holding for each other even when *evaluated even within each other's perspectives*. That is, we may hold that there's reason for you to do such and so when evaluated from my context α, from your context β, or even from any arbitrary context γ. Call this the *ontic* sense of normative universality: that we have reason to do something is ontically normatively universal when it holds no matter what the non-normative facts may be *when evaluated at any normative perspective*.[13]

[10] Or whatever your local "I'm not one of *those* people" supermarket is.

[11] Or whatever your local sensible supermarket is.

[12] You might think this rather partial and that morality ought to be impartial. But the strong forms of impartiality which treat all of our beliefs about the normative as on a par can't be held coherently with sensible subjectivism, as I pointed out in passing above. See Atiq (2016) for a more subjectivist friendly account of impartiality.

[13] This way of distinguishing two notions of reasons is similar to the two-dimensional interpretation of expressivism in (Peacocke 2003). (Schafer 2014a) develops a slightly more general view, *perspectivalism*, using the notion of assessor relativism.

To make this distinction slightly more precise, I'll introduce some machinery from Einheuser (2006).[14] Let a *context* be a pair $\langle c, n\rangle$ of a *circumstance* c–here a set of non-normative descriptive facts—and a system of norms n–here a set of conventions for what we have reason to do. A circumstance is a suitably large fragment of a possible world—large enough to support a group of agents with particular behavior indicating acceptance or rejection of a system of norms, systematic approval and disapproval of various actions, etc.[15] A context *gives rise* to a reason for an agent a to φ in a situation γ when the system of norms n, applied to the non-normative descriptive facts c, says that a has reason to φ in γ. We'll write this $\langle c, n\rangle \models R(a, \varphi, \gamma)$.[16]

We say that the system of norms n is *grounded* by a circumstance c when n is the system of pure reason facts (those which aren't dependent on particular non-normative facts about our perspective) arising from a circumstance c. What it is for a circumstance to give rise to a system of pure reason facts will depend on the correct view of metanormativity. Quite generally, we'll write n_c to indicate the system of norms grounded by c.

This way of proceeding explicates CONVENTIONALISM, but it doesn't presuppose it. If CONVENTIONALISM is right, then it's the behaviors and dispositions to accept various norms as holding (or deferral to a group of people who decide which norms hold, as in (Hart 1961)) that does the relevant work here by giving rise to a system of conventions more generally, which in turn gives rise, by means of CONVENTIONALISM, to a system of reasons facts. Then the actual circumstance $c_@$ grounds a system of norms n which says, when applied to the non-normative facts captured by $c_@$, that we have decisive reasons not to kick dogs since the conventional practices we accept (by means of our dispositions to accept governing principles which) rule out kicking dogs. More generally, given CONVENTIONALISM, our pure reasons will depend on our circumstances. If, alternatively, a non-subjectivist view is right, then there will be one set of *pure* reasons which will be grounded by any circumstance in some way. So our Einheuserian machinery makes room for,

[14]The development here is a bit more technical than the rest of the paper: the reader can skip it and work with the intuitive glosses if they're allergic to symbols.

[15]We use fragments of possible worlds to allow that different communities in the same world may have different normative behaviors. If this annoys, make the simplifying assumption that there is only one set of behaviors at any context and let a circumstance be a possible world.

[16]We will often simplify by assuming that a is part of c and that we don't have to deal with reasons arising for agents in context where they don't exist.

but doesn't presuppose, subjectivism.

We can vary our actual circumstance $c_@$ and our actual norms $n_{c_@}$ in two ways, thus obtaining two notions of universality.[17] What I've called *evaluative* universality can be defined easily enough for simple reasons claims:

$$\langle c, n\rangle \models \Box_e R(a, \gamma, \varphi) \text{ if and only if } \forall c'[a, \gamma \in c' \Rightarrow \langle c', n\rangle \models R(a, \gamma, \varphi)]$$

That is, a reasons claim $R(a, \gamma, \varphi)$ is evaluatively universal at a context $\langle c, n\rangle$ just in case a having reason to φ in situation γ—$R(a, \gamma, \varphi)$—holds at all contexts $\langle c', n\rangle$ which differ from $\langle c, n\rangle$ only in circumstance (i.e. non-normative facts) (and where a and γ still exist).[18] That is, when evaluating by the conventional norms n, no matter what circumstance we look at, a has reason to φ in γ.

Extending our definition in the obvious way, we get

$$\langle c, n\rangle \models \Box_e \varphi \text{ if and only if } \forall c' \langle c', n\rangle \models \varphi$$

for more complicated reasons claims with embedded occurrences of R. We can then find non-trivial evaluatively universal claims easily enough.[19] Consider "no one ever has a reason to be cruel merely for fun". Evaluated by a stringent version of our own moral outlook, this seems true no matter what the underlying non-normative facts are. Note that there is no entailment from subjectivism of the types described above to cases of evaluative contingency.

We'll have evaluative contingency only when n's verdicts depend on which context we're evaluating. Given CONVENTIONALISM, this will typically only be the case when we've got cultural relativist or, more plausibly, derivative "when in Rome" conventions. But many of our actual norms will manifestly not involve the idea that others have reason to do what their conventions suggest; in fact, we may actively disapprove of others following their conventions, etc. Anyways, regardless

[17] The use of $\Box$ below might suggest that I'm talking about two notions of necessity. My only objection to this terminology is that it brings in its wake more settled ideology than I'd like. However, if thinking in terms of necessity and contingency instead of universality and contingency helps you, feel free.

[18] I'll take the restriction that a and γ exists at c' as understood henceforth.

[19] There are complexities here that I'm ignoring; e.g. without restrictions on the range of φ, our definition marks various metaphysically necessary non-normative claims as evaluatively universal. See Woods and Maguire (2017) for related discussion.

of the contours of our actual conventions, subjectivism is by no means *committed* to evaluative contingency.

Defining *ontic* normative universality is also straightforward:

$$\langle c, n\rangle \models \Box_o\varphi \text{ if and only if } \forall c'\ \langle c', n_{c'}\rangle \models \varphi$$

That is, φ is ontically universal at a context $\langle c, n\rangle$ just in case that it holds in any context consisting of a circumstance c' and the norms grounded by c'. Presuming CONVENTIONALISM is true, when φ holds at the point of view of any context when the salient norms are those arising from the circumstance of that context. CONVENTIONALISM entails the existence of ontic contingency so long as different circumstances embody different systems of norms—as must be allowed to be the case on any plausible conventionalist view. Again, remember that the only plausibly universal reasons claims, in either sense, are pure reasons claims. Otherwise our reasons are typically held hostage to non-normative facts embodied in the circumstances of evaluation.

We can expand on our Einheuserian terminology to define two related senses of reasons. Let us say that a, in their context c_a, has an *ontic* reason to φ in γ just in case $\langle c_a, n_{c_a}\rangle \models R(a, \gamma, \varphi)$. And, fixing a set of norms n, we will say that a has an evaluative reason to φ in γ just in case $\langle c_a, n\rangle \models R(a, d, \varphi)$. In plainer language, a has an ontic reason to φ in γ just in case the norms governing *their* context give a reason to φ. And a has an evaluative reason, according to n, to φ in γ just in case n (applied to c_a) says a has reason to φ in γ.

It's now easy to see that a's reason to φ in c is ontically universal just in case a has an ontic reason to φ in γ in any context where a and γ exist and a's reason to φ is evaluatively universal just in case n yields a reason for a to φ in γ given any context c'. Having clarified the distinction between two types of normative universality and the two related sense of reasons, showing them obviously coherent, we turn to our central question about normative universality.

4. Normative Universality

As mentioned above, some theorists claim that treating normative facts as arbitrary or contingent would disrupt their functional role. Here's two sources, one old and one new:

> if one does not wish to deprive the concept of morality of all truth and all relation to any possible object whatsoever, then one cannot dispute that its law is so extended in significance as to be valid not merely for human beings but for all reasonable beings whatsoever, and not merely under accidental conditions and with exceptions but with absolute necessity... (Kant 2002, 24)

> ...if normative reasons were indeed relative, then mere reflection on that fact would suffice to undermine their normative significance. (Smith 1994, 172)

There are two immediate responses to this. First, one borrowed: Kant seems to be articulating a brute intuition about the ontic universality of normativity (or, at least, the moral bit of normativity): It's *categorical* and *binding* on all rational individuals. But this brute intuition is hard to defend. As Foot (1972) showed, the categorical form of (moral) normative judgments is insufficient to validate Kant's intuition. Obviously contingent etiquette facts likewise have categorical form; it's not "respect your elders if you don't want to be rude", but "respect your elders!"[20] And it seems equally weird to respond to the politeness imperative with "but I don't care about politeness" as it does to respond to the moral imperative with "but I don't care about morality".[21] But what then justifies the claim that normative judgments are ontically universal?

Now, one blue: Difference in conventional morality is simply a fact of life and one non-theory driven ordinary moralizers have learned to live with. Diversity in prudential opinion is likewise fairly common; risk aversion comes in different levels, as do views about what a good life consists in and how to best care for ourselves. Even diversity in epistemic norms is rife and sometimes intuitive (just think of various inductive rules and standards of statistical significance we might adopt). Reflecting on this fact doesn't undermine the normative significance of *our views* of risk, morality, and epistemology *for us*; far from it.

[20] Miss Manners, treats this as evidence that etiquette facts are moral facts (Martin and Stent 1990). This, though, conflicts with the strong intuition that etiquette facts have only local force, though it's fair to say that many etiquette facts make overlap, in their content, with moral facts—don't offend! And many moral facts have as part of their content *respect*, which in turn implies we ought to pay attention to etiquette. But the claim that *all* etiquette facts are moral is too much to bear.

[21] See Woods (forthcoming, §10.3.1) for an explanation of why this is.

Consider that the seemingly undeniable tendency to treat (at least some of) our aesthetic judgments, especially judgments of mere taste, as sourced in our own standpoints doesn't undermine our application of these standards to ourselves. If we understand Kant and Smith as saying that accepting relativism undermines taking all normativity seriously, their claim is implausible. When we take a standard seriously, it plays a role *for us* in regulating *our* behavior, contingent or no.[22] Maybe the claim is that it would be better if we believed our normative reasons to be non-relative. But this is not at all obvious. Often the fact that my reasons would change if I acted thusly seems intuitively explanatory of why I shouldn't act thusly; this not only is consistent with subjectivism, but seems to nearly presuppose it.

A better thought is that normative facts hold not only for ourselves, but also for others, regardless of whether they share our normative beliefs or not. Further, the thought continues, this is part and parcel of why we engage in normative theorizing to begin with. Perhaps we should understand Kant and Smith as claiming that the functional role of normative theorizing requires that we be able to criticize and correct the actions of others. Lenman, discussing Smith, glosses what this might mean:

> If we came to see our moral commitments in this way we'd rightly panic because we'd no longer be able to take seriously the idea of disapproving of someone for failing to share them. (Lenman 1999, 166)

That a central part of the functional role of normative judgment includes disagreement and evaluation of others seems true.[23] It'd be difficult to come to grips with a notion of normativity that didn't allow us to bring our normative beliefs to bear on each other in attempting to guide not only our own actions, but also the actions of others.

[22] We can treat taking a formal standard seriously, roughly, in terms of us taking being out of step with it as non-instrumentally undesirable. See Woods (2016, forthcoming) for details and complications.

[23] I won't address the full literature on disagreement in ethics and aesthetics. Since disagreement is just one of several properties I think the subjectivists can accommodate with evaluative reasons, I just sketch one way disagreement is possible. See Finlay (2014) and Ridge (2014, chapter 6) for further accounts of disagreement for subjectivist and expressivist views which could be adapted to my purpose here.

We should be careful here though. Some normative standards clearly don't work this way. My judgment that it's rude to expose my feet to others gives me no reason to keep my shoes on in a Turkish household; their sense of it being rude to object in class gives them no reason to behave deferentially in *my* classroom. Quite the contrary. Different standards of etiquette apply in each case and we are well aware that our behavior should change accordingly, at least when we think carefully about it. That this isn't the case with morality may be due to differences in the functional role of our normative and our etiquette judgments.

So perhaps this is the essence of Kant and Smith's point: we view normative reasons as holding for others as well as ourselves; our evidence for this is our willingness to bring such considerations to bear in our discussions with one another about how we ought to act; and the explanation of this may, perhaps, be found in the functional role of normative theorizing. Other potential aspects of the functional role of normative theorizing includes evaluation of the reasons of others, contentful disagreement about what to do, planing for what to do in hypothetical situations, and evaluation of the reasons we'd have there. Reasonable subjectivist views should capture these aspects of normative judgement's functional role. As I'll show in the remainder of this essay, capturing these aspects does not require the ontic universality of normativity—something far weaker will suffice. We will now turn to the first aspect—evaluating the reasons of others.

5. Subjectivism and Criticizing Others

Consider again a particular form of aesthetic judgment, judgments of taste. It's commonly (and correctly) believed that tastes are relatively subjective. I am no fan of chocolate; it's bearable, but generally unremarkable. This is strange to most people. Assessed from their gustatory standpoint, I have reasons to eat more chocolate. "Assessed from their gustatory standpoint" here means evaluating what gustatory reasons I have according to their tastes. The evaluative reasons I have, from their standpoint, will differ from the evaluative reasons I have according to my own gustatory standpoint. Plausibly, I've evaluative reason to eat chocolate according to the average individual, but no ontic reason at all.

Consider likewise folks driving on the wrong side of the road. This is a joke—kind of. The funny thing is that we *really do* feel like other folk drive on the wrong side of the road, even though we know quite well that it's entirely arbitrary which side of the road a community drives on. The intuitive force of the approval

of right-driving over left-driving remains, even when we know quite well that the choice of side is arbitrary. Of course, this is quite irrational *in a sense*. We all know full well that tastes differ between individuals and that my tastes give you no reason to eat what I would eat were I to be in your shoes, but with my tongue.[24]

In other cases where we have explicitly subjective reactions, recognition of this fact undermines claims to correctness and incorrectness.

> ...once we bring other perceptual systems into view, then provided they are equally discriminatory, we lose any very robust attachment to the idea that ours is right and theirs is wrong. ... People who taste phenolthiourea the other way are not wrong, just different. But there is no reason to suppose that this ambivalence extends similarly to the case of value. (Blackburn 2006)

But I submit that many of us feel a drive to criticize the gustatory sensibility of others *even though* we know full well that tastes are subjective—or, anyways, so my experiences being a chocolate deviant suggest. Our actual gustatory standpoints seem to include the idea that we are in a position to recommend someone eat against what we know their tastes to be, to criticize their tastes for being bizarre, and so on. And, *in a sense*, this is quite rational as tastes are plastic and, from your point of view, chocolate really is delicious. Continued experience with a particular thing, be it chocolate or wine, tends to breed some taste for it, which from your perspective is a valuable thing.[25] Moreover, general convergence in our tastes is useful and desirable—when deciding where to go for dinner. And we might want ourselves, if we were to find ourselves in a position like mine, to go through the pain of developing an affection for chocolate, at least as we *now* view that tragic possibility.

The same explanation cannot be given for ontic aesthetic reasons. It's irrational *tout court*, not merely irrational in a sense, to think that we all have gustatory reasons to eat chocolate when assessed from within our own gustatory perspective.

[24] Alex King suggests (personal communication) that I'm staking a lot on the subjective nature of culinary norms and that perhaps the general aesthetic case, including more objective aspects of the aesthetics of food, is different. The issue is difficult, but I only really need the gustatory case for my point that we *can* use what seem obviously subjective normative standards to fill the functional role that we need for putatively non-subjective ones like morality. Thanks also to her for more general discussion of these issues.

[25] The general aesthetic case is more complicated, as can be seen by the results of Meskin et al (2013).

Some of us—ahem—do not. There is no gustatory reason for me to pursue the modification of my desires that continued exposure to chocolate might provide.[26] Typically, my only reasons to develop a taste for chocolate are prudential reasons like not standing out or health benefits.

Yet this type of aesthetic irrationality is a pervasive feature of ordinary aesthetic discourse. Someone who differs significantly enough from you aesthetically is often thought to simply be mistaken about their own tastes, as if such a thing were generally possible. But this is projection of some kind of mistaken universality on a bit of our experience we know, on reflection, to be non-universal. Thus the *explanation* of why we might hold onto ontic aesthetic universality doesn't justify holding on to it, *unlike* the case of evaluative aesthetic universality where aesthetic norms evaluating others tastes negatively are coherent and socially useful.

As with aesthetic reasons, so with normative reasons. We can make sense of evaluative normative reasons and thereby evaluative normative universality. And the reasons given above for evaluating the aesthetic reactions of others are even more compelling in the normative case. If we are often willing to claim aesthetic reasons hold for someone even when they clearly aren't reachable from within their aesthetic standpoint, even given that we know in our hearts how subjective tastes are, it seems very plausible that we are willing to claim normative reasons hold for someone else when they aren't clearly reachable from within their normative standpoint. Our normative standpoint need accept no restriction on our evaluation of other's actions, even when they clearly don't share our norms.

Applying *our* normative notions to *your* actions yields a definite and sensible verdict—even if it's not obviously a verdict you should share.[27] It's here that we find the truth in Smith's idea that viewing our own perspectives as arbitrary would undermine them. If we viewed our reasons as holding *evaluatively* only for us or those with nigh-identical norms, then we'd be hard pressed to make sense of the role of normative reasons in evaluating others; a role which is clearly part of our actual normative practice. Luckily, since exertion of social pressure by means of explicit normative pronouncement is reasonably effective as a tool of social coordination, even of those with divergent norms, there's good reason to evaluate

[26] Of course, there might be "taste-directed" non-gustatory reasons (akin to "truth-directed" non-epistemic reasons) to do so since I might get to enjoy more tastes I otherwise enjoy were I to love chocolate more.

[27] Consider the similar discussion of defined notions of evaluation in (Schafer 2014a).

others.[28]

I've argued that we can reasonably evaluate others actions by means of our own norms even when we know that our norms differ by analogizing it with a clearly subjective set of evaluative judgments: aesthetic judgments. But there remains the question of how to understand our frequent engagement in explicit argument with others about normative matters and our attempts to convince them to modify their normative standpoints. Evaluation is one thing; reasonable hashing out of normative difference another.

If someone's normative or aesthetic standpoint is significantly different from ours, this practice might seem pointless. *De gustibus non est disputandum*, after all. Yet, we dispute about taste all the time.[29] So we have to explain why would we do such a thing if our normative standpoints aren't ontically universal. We turn to this now.

6. How Subjectivists Should Disagree

In order to address this question, it's useful to set aside a class of cases where normative disagreement *doesn't* make sense. If someone has a coherent and deep commitment to a normative standpoint *radically* different from ours, then dispute can be pointless. Of course, as mentioned above, this does not mean that we can't evaluate their actions according to our own normative standpoints and that, in an important sense, it's totally reasonable to do so. But actually locking horns with them about their perverted norms doesn't really make sense.[30]

This echoes Ayer's (1946) neglected discussion of moral argumentation. He analyzes moral argumentation as proposing various "deeper" commitments until our interlocutor latches onto one of them. Punching him is okay?; Well, do you think hurting people is fine?; Oh, do you think causing pain unnecessarily is fine? and so on. Ayer suggests that if this procedure does not work, or if we don't uncover false non-normative beliefs in the process, we tend to abandon the argument. Sim-

[28](Manne 2014) also suggests similar justification for persisting in offering reasons to others even if our offerings won't have uptake for them.

[29]I've always preferred the punny rendition of the phrase in English anyway. There's no *accounting* for taste is a nice way to express that one searches in vain for why someone likes lettuce.

[30]Compare the discussion of the pointlessness of discussion with "the others" in (Lenman 2014, 243).

ilarly, if we can see in advance it won't work, then it seems unreasonable to start arguing with someone at all.[31]

Ayer neglects the role of argument in *irrational* change of belief—we can bring people to agreement by applying argumentative peer pressure. Sometimes this is even reasonable.[32] But, when someone has suitably robust views and a suitably argumentative nature, it may be pointless to disagree:

> Suppose...I am trying to convince a man who is nasty to his wife to treat her more nicely, or with more consideration. In this endeavor, I repeatedly press my concerns on him, and in a variety of ways. Finally, he says to me—borrowing Williams' wording here—"I don't care. Don't you understand? I really do not care." That is, he doesn't care directly about treating his wife more nicely. Nor does he care about any of the goods which would be promoted or instantiated by so doing....
>
> Here's the intuition I have, and want to invite you to share, now: when we learn that this man cannot be motivated to lift his game merely by continuing to carry on with the conversation, something has now changed in the normative and dialogical space between us. (Manne 2014, 102–103)

Manne's point is that when there's really no common ground for rational conversation, the only reason remaining to argue is browbeating one's interlocutor into moral compliance.[33] Our normative standpoints differ to simply too large a degree.

Thankfully, such cases are relatively rare; we can put them aside. In better cases when our interlocutor is reasonable—to borrow Manne's expression, when they are open to 'rational' discussion—then we presume overlap in our beliefs about what considerations normatively matter. It's plausible that such overlap is required to make sense of rational discussion at all.[34]

---

[31] Though, as Jimmy Lenman usefully reminded me, the situation is more complicated in practice. We might, for instance, be *performing* our argument for an audience.

[32] See also Street (2009) and Sobel (2014) for this point.

[33] Manne is concerned with cases where motivational internalism fails, but the point generalizes.

[34] See Lenman (2014) for related arguments to the effect that the whole point of normative communication — at least in typical and important cases — is to come to convergence in our

Carballo and Santorio (2016), for example, argue that for rational communication about normative matters to have a point, both participants must take as part of the presupposed common ground that:

- there are potential norms we might accept—where such norms are either derivative or fundamental—which are not ruled out in advance of each other by either party
- the participants ought to come to some convergence.[35]

These conditions are generalizations of the typical Stalnakarian account of communication, so we needn't even invoke anything special in understanding normative communication. We can thus understand communication about normative matters in terms of coming to agree on norms or attitudes we can share and usefully apply in planning our actions and evaluating the world around us.[36] If this is right—and I suspect it is—then we should assume that our typical interlocutors share some part of our normative standpoint or, at a minimum, a set of commitments about how to go about improving our norms.[37]

Restricting our norms to these overlapping sections of our normative standpoints, there is a reason for someone to do something by my lights (there's ontic reason to do something) if and only if there is a reason for them to do it by yours. By arguing about normative matters with the aim to coordinate, we are sometimes putting forth factual claims about what reasons exist in the shared background and sometimes proposing to extend this shared background by adding norms to

---

normative standpoints. As Lenman rightly stresses, this point is easily satisfied in principle, even if far too infrequently in practice.

[35]They go on to argue that both participants should agree that there is a unique convergence point, but this is overly strong for subjectivist views. If this were true, however, it would just make my point all the stronger.

[36]Finlay (2009 §7.1) argues compellingly that participants to such conversations often *do* presuppose overlap in normative outlook (in his terms, shared ends). This strengthens the point that normative argumentation does and should proceed in a sensible way even if subjectivism is true. He goes on to suggest—as I did above—that even when such presuppositions are false, they're useful to maintain. Williams (1995) similarly glosses the phenomenon in terms of launching our normative assertions into the mass of humanity with the optimistic hope that they'll find purchase in our interlocutors. See also Lenman (2014) for further discussion.

[37]Of course, this latter is simply another part of a normative standpoint, but it's useful to highlight that we can get on with normative communication if we either share norms or certain "metanorms" about how to revise norms (for instance, by sharing a common conception of what moral evidence consists in) . Thanks to Pekka Väyrynen for discussion.

the common ground and making the appropriate modifications to one's overall normative standpoint. Since the shared background will typically be a mere fragment of our overall normative standpoint, which may or may not be complete, accepting the proposal of our interlocutor will often involve retracting or modifying some fragment of our normative standpoint.

Here are two examples. Suppose neither you nor I take a stand on whether third-trimester abortions are permissible, but we agree (a) about many of the non-normative details, (b) that pain and suffering is bad, we've reasons to avoid taking life, etc, and (c) that we've reason to make our view more coherent and natural when possible. You propose that we've reasons to avoid third-trimester abortions on grounds that fetuses are well-developed by that point and the likelihood of suffering is non-negligible. I accept this claim, add it into the common ground, and thereby expand the overlap of our normative standpoints.

Now, suppose that the case is as above, but I think we've reasons to avoid third-trimester abortions and you disagree. I propose as before, suggesting that if you accept b, you *should*—on abductive grounds—accept that we've reasons to avoid third-trimester abortions. You come to accept that we have reasons to avoid third-trimester abortions, retracting your earlier view.

If *our* communication about normative matters looks like this, then we do not need ontic or even evaluative universality to make sense of normative communication, or anyways, the normative communication that actually matters. All we need is presupposition of significant overlap in our normative standpoints.

Presuppositions may be false; there may be less overlap in our normative standpoints than we hope. But the more atypical the case, the more pointless it will seem to engage in rational discussion instead of engaging in browbeating, the exertion of social and peer pressure, and the like. Harkening back to the opening remarks from Foot, Street, and Stroud, those who share *none* of our normative standpoint[38] seem inhuman and alien, unlike our typical conversational partners. Since all we need in the presumption of overlap (on both parts), and this follows

[38]As above, they must also fail to share common principles for improvement of sensibilities (where there's common views of evidence for what constitutes improvement), beliefs about the functional role of normative judgment, etc. This hedge being noted, the idea is clear enough. To be honest, I think we reach incomprehension easier than this, but I want to be ecumenical or, at least, as ecumenical as I can stomach.

from viewing our conversational partners as relatively like us, there's good reason to think that normative communication, even for subjectivists, is obviously possible and plausibly pervasive. We turn now to addressing a related aspect of the functional role of normative judgment: hypothetical planning.

7. Subjectivism and Hypothetical Planning

We often consider hypothetical and counterfactual situations to prospectively *plan* what we should do—to test and hone in our actual reasons (Gibbard 1990). We might want ontic universality here since we want our garnered knowledge to bear on analogous circumstances we might actually face. We want both that the reasons on which we'd then act would be genuine reasons for action as well as our testing of general reasons would be guaranteed by ontic universality.

As Schafer (2014) notes, we use this kind of reasoning all the time:

> Nothing is more natural or more common than for us to consider someone's situation and to form some hypothetical plan for action for the situation in question. For example, suppose I... find myself wondering what I "would have done" had I been in Napoleon's shoes. When I ask myself this question, I am not asking myself a descriptive question concerning what someone with my psychology would have done in that situation. Rather, I am asking myself what to do in such a situation. In other words, I am forming a hypothetical plan for action...(Schafer 2014b)

But there is a worry here: why bother forming a plan for what I would do in Napoleon's shoes as I'll never be there, and I'll never share his norms?[39] Even less the norms of Caligula or Elizabeth Báthory. What role do these bizarre hypothetical plans play in my actual plans for situations I'll encounter? Perhaps our pure reasons should be treated as ontically universal if this kind of hypothetical planning need inform our actual decisions; but since many of the problematic situations seem to do no such thing, it's hard to see why we should require ontic universality.

[39] See also Finlay (2014 §8.6) on weird consequences of Gibbard's account of hypothetical planning.

Of course, it *can* be reasonable to evaluate Napoleon's situation by *our norms*. To be sure, it would make more sense for a general to engage in an evaluative judgment about what Napoleon should do than it would for most of us, but the point is the same. It's reasonable to evaluate what *we*—people with our normative sensibilities—would do in similar situations. This doesn't give us much reason to think that ontic evaluation of Napoleon's situation is anything more than a curiosity. What would the fact that Caligula has ontic reason to plan a massacre tell us about what we'd have reason to do in a similar situation?

The regulative role of considering ontic reasons in bizarre abstract hypothetical situations is thus much less important that it initially seemed.[40] Moreover, it's doubtful that the anti-subjectivist intuitions about such situations usually bandied about are clearly probative. Imagining the *full* normative outlook of Parfit's Future Tuesday Indifferent, for example, (someone who only prospectively cares only about non-Tuesday pain) seems nigh impossible, so our intuitions aren't clearly probative—plausibly we're importing our norms illicitly in claiming they have reason to avoid pain on future Tuesdays. Obviously such a situation is then irrational, but it's irrational since there's internal conflict between the norms and desires we're illicitly importing and those part of the stipulated example.

This point should remind us again of Wittgenstein's Woodsellers and how difficult it is to understand their way of counting, especially when we've only considered how they measure wood. The point is not that we cannot make out how they measure wood; it's understanding how it integrates into an overall system of measuring and the role of counting in various aspects of our lives. Measuring wood by area is tantamount to throwing out everything we know about measuring. It's entirely alien to us. As with wood, so with reasons.

Why, then, is it so natural and easy for us to pass seemingly ontic judgment in such counteractual cases? Presumably, it's because our judgments about what we ought to do in Caligula's shoes are informed by the gross majority of *our* normative sensibilities *and* we assume that Caligula, different as he is, nevertheless shares many of our norms. Again, it's very plausible that there is broad and per-

[40]This is true for morality and normativity generally, but not clearly true for other "when in Rome" systems of norms like those of fashion, aesthetics, and etiquette. There it seems much more important to consider what our reasons would be in that context instead of what we think, from outside, the reasons they have are. Still, and crucially, we don't need ontic universality for this. Thanks to Catharine Diehl for discussion.

vasive overlap between any *conceivable* normative outlook. Maybe it seems okay to cross against the light and shoplift from large chain stores to me and not to you, but random stabbings and complete and explicit disregard for others[41] are *out*. And asking ourselves the question of what we'd do in Caligula's shoes plausibly presupposes this.

Something like this underlies Foot's famous observation:

> But it's interesting that the people of Leningrad were not similarly struck by the thought that only the contingent fact that other citizens shared their loyalty and devotion to the city stood between them and the Germans during the terrible years of the siege. Perhaps we should be less troubled than we are by fear of defection from the moral cause…

They were not struck by this exactly because it would be relatively *inhuman* to have no loyalty and devotion to one's city, family, and friends; it seems literally inconceivable that this could be a pervasive feature of our fellows. We accept that not everyone shares our sensibilities, but looking through the eyes of those who do not seems rather difficult. Again, where even to start?

We might thus mistake Caligula's evaluative reasons (roughly corresponding to what *we* would do in similar circumstances) for *ontic* reasons. Given the presupposition that we enjoy broad overlap in sensibility and normative outlook, the questions come to largely the same.[42] So it's no surprise that we'd mistake the evaluative reasons Caligula has to avoid massacre-planning for ontic reasons. There's thus not much cause for worry that subjectivism delivers the wrong verdict on these situations;[43] such intuitions are plausibly based on illicit readings of the cases.

If something more precise is desired, we could put the upshot this way. We need only *Human universality* (writing $c \sim_h d$ for d a context similar to ours in the conventions, psychologies, and behaviors of the agents within):

$\langle c, n\rangle \models \Box_h R(a, \gamma, \varphi)$ if and only if $\forall c'[a, \gamma \in c' \wedge c' \sim_h c \Rightarrow \langle c', n_{c'}\rangle \models R(a, \gamma, \varphi)]$

[41] See Street's (2005) discussion of the rational social insect.

[42] Note that this means that evaluative reasons are enough even for non-subjectivist accounts of hypothetical planning, so we haven't lost anything here by our subjectivism.

[43] I won't explicitly list the many heralds of this mistake, though the reader won't struggle hard to find them.

for normative judgments to play the role in organizing our actual behavior that Gibbard and Schafer suggest. That is, we only need that our reasons are the same in any *human* contexts, contexts more or less similar to ours.[44] And no reasonable subjectivist position needs to force us to believe normativity is not humanly universal. Intuitions to the contrary can be explained away by a combination of Foot's observation about humanity and Stroud's observation about perverse method of underexplicated cases. We now turn discussing one last aspect of normative judgment, one which generalizes Smith and Kant's point from above and continues the point that our intuitions about counter-normative contexts aren't probative.

8. Counternormative Counterfactuals

The following counternormative conditional seems false:

> If local moral conventions demanded random cruelty, I would have strong reason to be cruel

We seem not to have strong reasons for being randomly cruel, whether or not our local conventions endorse it.[45] But, since we have distinguished two senses of universality, we should also distinguish two senses of these conditionals.[46] We can distinguish the evaluative normative counterfactual conditional:

$\langle c, n\rangle \models \varphi \mathrel{\square\!\!\rightarrow_e} \psi$ iff in the closest circumstance c' in which $\langle c', n\rangle \models \varphi$, $\langle c', n\rangle \models \psi$

from the ontic normative counterfactual conditional

$\langle c, n\rangle \models \varphi \mathrel{\square\!\!\rightarrow_o} \psi$ iff in the closest circumstance c' in which $\langle c', n_{c'}\rangle \models \varphi$, $\langle c', n\rangle \models \psi$

So long as we do not have as a part of our normative standpoint that the conventions of people matter for whether they have reasons for what they do—and most subjectivist views endorse no such thing—then the closest contexts in which people approve of random cruelty is one where random cruelty is wrong. So the

[44]Note that this is a restriction of *ontic*, not *evaluative*, universality. We can still have full bore evaluative universality in tandem with ontic universality; in fact, I think this is more or less the usual situation for a suitable precisification of 'human'.

[45]Modulo lightweight reasons that arise from distinct contexts, like promising to be cruel. I am committed elsewhere to the view that such reasons exist, but are nearly always outweighed. See Woods (2016, forthcoming).

[46]I work with a Lewis-style "variably-strict" analysis of counterfactuals here for simplicity (Lewis 1973). My point can be generalized across a range of analyses of counterfactuals.

evaluative counternormative is false.

Is this enough to downplay the intuition that such conditionals are intuitively false? It ought to be. It's entirely unclear why we'd think that the ontic counternormative is obviously false. But I fear that this will be insufficient to convince most people, so let me offer another consideration. Many of ontic counternormatives are also false *even given subjectivism*. Consider the closest possible world in which kicking dogs is endorsed by our normative standpoint, but where we don't have massively false beliefs about animals.[47] What would such a world be like? Our disapproval of kicking dogs is grounded in our empathy towards animals, the intimate role dogs play in our lives, that kicking them causes pain, etc. This sits badly with approval of kicking dogs. So the nearest world in which kicking dogs is smiled upon is one falsifying one of these facts. Perhaps dogs there enjoy being kicked. In that context, it's not clear that kicking dogs is bad.[48]

Suppose, now, that we do not feel any empathy towards animals. We have two types of cases. First, we might leave the rest of our psychology alone. But given the resulting massive incoherence in our psychologies, such a world will be very different from ours. Sympathetic reasons to avoid randomly cruelty, after all, are more or less normative bedrock. Given this, the above point about dog-kicking applies. A world where we endorse dog-kicking with no deeper explanation is so far away that it's not relevant for evaluation of the counternormative.[49]

Suppose instead that we make the requisite changes in the rest of our psychology. Then, as in the last section, I submit that our intuitive judgments about ontic reasons aren't probative. Imagine what it would be like to have a psychology approving of something like serious random cruelty without assuming some massively mistaken set of beliefs about the world. I'm guessing that you can't; I definitely cannot. This, I believe, underlies why it's so difficult to even portray a seriously deranged character (say, in a show like Dexter); the tendency to "humanize" them

[47] In general, we need to be careful here about instrumental norms and information. For example, if our anti-dog-kicking norm is grounded in more basic norm of avoiding causing pain to feeling creatures, then false beliefs about dogs being feeling creatures will explain why we endorse dog-kicking. Still, we would have ontic reason to avoid kicking dogs since they, in fact, are feeling creatures.

[48] For more discussion of counterconventional conditionals, see Einheuser (2006).

[49] Of course, you can freeze this condition into the prejacent of the counternormative, but evaluating such complex prejacent-ed conditionals is extremely fraught.

is nearly immediate.

So there are three interlocking responses to intuitive "false" judgments about counternormative conditionals. First, we might be mistaking our evaluative reaction for an ontic one. Evaluatively, our reaction is totally apt. Second, we might be mistaking ontic reaction for one case for an ontic reaction to a different case—the case of a group differing from us in this particular, but agreeing on nearly every other part of our normative standpoint. But such cases are too remote to play a role in the truth or falsity of the counternormative. Finally, we might be treating our reaction to the case just described as our reaction to the case of someone with a coherent, but inhuman, psychology. But such intuitions aren't probative, as I've argued above.

So counternormatives pose no real problem. The relevant evaluative counternormatives are typically false, and, for ontic counternormatives, once we've disentangled what they really are, we can see that our intuitive reactions aren't probative. Given this, we cannot assume that they're false without begging questions against the subjectivist.

9. Conclusion

I have distinguished and defined two separate notions of normative universality—ontic and evaluative universality—and argued that even though subjectivism denies ontic universality, this does not undermine normative judgment playing its usual functional role. Evaluative universality suffices to make sense of evaluating the reasons of others, significant overlap in our reasons suffices to make sense of normative disagreement and evaluating hypothetical cases, and the resulting types of counterfactuals are unproblematic—the relevant evaluative counternormative counterfactuals are false, as should be expected, ontic counternormative counterfactuals are either false, unworrisome, or our intuitions about them are not probative.

At this point, it might be worried that there are real life cases of repulsive moral behavior. All too true. But charity demands that we should try as hard as possible to see the culprits as not making moral errors, but factual ones. When we cannot find such errors, we should look for swamping beliefs—such as religious, cultural, or capitalist convictions that cut against human empathy—or mistakes of

reason arising from mistaken application of beliefs to norms.[50] And when both of these fail, only then should we look to see what sort of affective or moral difference would explain why we and they have such differing normative outlooks. If this procedure is reasonable, as I hope and pray it is, then plausibly we generally presuppose that others are relatively like us.

When we cannot so explain grotesque behavior, it seems to me that we are in a position to these folks as we are to people who approve of random cruelty on alternating Wednesdays. It's very difficult to understand how someone could willingly throw adulterers off of high towers without presuming that they believe something that's simply false. This normative outlook is not just different from ours; it's more or less incomprehensible. Or, anyways, so I find it. And, judging by how often people react to immoral behavior with expressions like "I simply don't understand how you could think that", I and my fellow subjectivists are not alone in this.

Even when we can make sense of evaluating the reasons of someone from a differing perspective, it by no means follows that we *should* take such evaluations seriously or afford them the role that our assessments of reasons typically plays in attributions of punitive actions like blame and criticism. That's a non-obvious normative claim. Likewise, it's not obvious that we should revise our tendencies to accept norms demanding evaluative normative universality, just as it's not obvious that we should revise our tendencies to accept norms demanding evaluative aesthetic universality. These tendencies play a useful prudential role in facilitating coordination, one which would be hampered by being overly tolerant of the viewpoints of others.

On balance, the assumption that normative facts are ontically universal is on shaky footing. If there's an argument here, it won't come from subjectivism interfering with the functional role of normative judgment. Without begging questions against a subjectivist point of view, it's unlikely to come from anywhere. This is not to say that there aren't other reasons to reject subjectivism. But the putative ontic universality of normative facts is not one of them. This fact seems to me to significantly strengthen the already quite considerable case in favor of subjectivist views of normativity.

[50] See Sobel (2014) for useful discussion of the role of mistaken beliefs in explicating seemingly immoral and amoral behaviors.

References

Atiq, Emad. 2016. 'How to be Impartial as a Subjectivist'. *Philosophical Studies 173*(3): 757–779.

Ayer, A. J. 1946. Language, Truth and Logic. *Victor Gollancz Publishers*.

Blackburn, Simon. 1998. *Ruling Passions*. Oxford University Press.

Blackburn, Simon. 2006. 'Must We Weep for Sentimentalism?'. in *Contemporary Debates in Moral Theory* ed. J. Dreier: 144–159. Wiley Publishers.

Carballo, Alejandro. Perez. and Paolo. Santorio. 2016. 'Communication for Expressivists'. *Ethics 126*(3): 607–635.

Dreier, Jamie. 1990. 'Internalism and Speaker Relativism'. *Ethics 101*(1): 6–26.

Dworkin, Ronald. 1996. 'Objectivity and Truth: You'd better believe it'. *Philosophy & Public Affairs 25*(2): 87–139.

Einheuser, Iris. 2006. 'Counterconventional Conditionals'. *Philosophical Studies 127*(3): 459–482.

Finlay, Steve. 2014. *Confusion of Tongues: A Theory of Normative Language*. Oxford University Press.

Foot, Phillipa. 1958. 'Moral Beliefs'. *Proceedings of the Aristotelian Society 59*: 83–104.

Foot, Philippa. 1972. 'Morality as a System of Hypothetical Imperatives'. *The Philosophical Review 81*(3): 305–316.

Harman, Gilbert. 1975. 'Moral Relativism Defended'. *The Philosophical Review 84*(1): 3–22.

Gibbard, Allan. 1990. *Wise Choices, Apt Feelings: A Theory of Normative Judgment*. Harvard University Press.

Hart, H.L.A. 1961. *The Concept of Law*. Claredon Press.

Kant, Immanuel. 2002. *Groundwork for the Metaphysics of Morals*. Yale University Press.

Lenman, James. 2014. 'Gibbardian Humility: Moral Fallibility and Moral Smugness'. *Journal of Value Inquiry 48*(2): 235–245.

Lenman, James. 1999. 'Michael Smith and the Daleks: Reason, Morality, and Contingency'. *Utilitas 11*(2): 164–177.

Lewis, David. 1989. 'Dispositional Theories of Value II'. *Proceedings of the Aristotelian Society 63*: 89–174.

Lewis, David. 1973. *Counterfactuals*. Blackwell Press.

Manne, Kate. 2014. 'Internalism about Reasons: Sad but True?'. *Philosophical Studies 167*(1): 89–117.

Martin, Judith and G. S. Stent. 1990. 'I think; therefore I thank: A philosophy of etiquette'. *The American Scholar 59*(2): 237–254.

Meskin, Aaron and Mark Phelan and Margaret Moore and Matthew Kieran. 2013. 'Mere Exposure to Bad Art'. *British Journal of Aesthetics* 53:(2): 139–164.

Parfit, Derek. 1984. *Reasons and Persons*. Oxford University Press.

Peacocke, Christopher. 2003. *The Realm of Reason*. Oxford University Press.

Ridge, Michael. 2014. *Impassioned Belief*. Oxford University Press.

Schafer, Karl. 2014a. 'Constructivism and Three Forms of Perspective-Dependence in Metaethics'. *Philosophy and Phenomenological Research 89*(1): 68–101.

Schafer, Karl. 2014b. 'Doxastic Planning and Epistemic Internalism'. *Synthese 191*(12): 2571–2591.

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

Smith, Michael. 1999. 'The Non-arbitrariness of Reasons: Reply to Lenman'. *Utilitas 11*(02): 178–193.

Sobel, David. 2009. 'Subjectivism and Idealization'. *Ethics 119*(2): 336–352.

Sobel, David. forthcoming. 'Subjectivism and Reasons to be Moral'. in *From Valuing to Value*: 16-42. Oxford University Press.

Street, Sharon. 2005. 'A Darwinian Dilemma for Realist Theories of Value'. *Philosophical Studies 127*(1): 109–166.

Street, Sharon. 2009. 'In Defense of Future Tuesday Indifference: Ideally Coherent Eccentrics and the Contingency of What Matters'. *Philosophical Issues 19*(1): 273–298.

Street, Sharon. 2010. 'What is Constructivism in Ethics and Metaethics?'. *Philosophy Compass 5*(5): 363–384.

Street, Sharon. 2012. 'Coming to terms with contingency: Humean constructivism about practical reason'. in *Constructivism in Practical Philosophy*, eds. J. Lenman, James, and Y. Shemmer: 40–59. Oxford University Press.

Stroud, Barry. 1965. 'Wittgenstein and Logical Necessity'. *Philosophical Review* 74:(3): 504–518.

Wiggins, David. 1998. 'A Sensible Subjectivism?'. in *Needs, Facts, and Values* ed. D. Wiggins: 185–215. Claredon Press.

Williams, Bernard. 1995. *Making Sense of Humanity*. Cambridge University Press.

Wong, David. *Moral Relativity*. University of California Press.

Woods, Jack. 2016. 'The Normative Force of Promising'. *Oxford Studies in Normative Ethics* (6): 77–101.

Woods, Jack. and Barry Maguire. 2017. 'Model Theory, Hume's Dictum, and the Priority of Ethical Theory'. *Ergo* 4:(14): 419–440.

Woods, Jack. forthcoming. 'The Authority of Formality'. *Oxford Studies in Metaethics* (13).

Velleman, J. David. 2013. *Foundations for Moral Relativism*. Open Book Publishers.