

Footing the Cost

(Of Normative Subjectivism)*

Jack Woods
j.e.woods@gmail.com

EXTREME Draft: do not cite without permission though
comments very welcome.

1. Introduction

Briefly scanning recent writing on normativity, we might well say that in (metaethical) myth and legend, the cost of countenancing the possibility of fundamental rational divergence in our normative outlooks is high—pretty near the end of the world. But this terrible cost is far more frequently assumed than argued for and, on inspection, it is difficult to see any non-question-begging arguments that recognizing the possibility of rational normative divergence should have any rational impact on our normative outlooks. I want argue even if difference in normative outlook underwrites differences in normative reasons, it is not clear that reflection on the possibility of such divergence in our normative reasons demands any substantial changes in *our* judgments about our actual normative reasons.

My argument for this is abductive; I will articulate a number of considerations that might push us towards thinking normativity should be necessary and argue that none of them requires necessity of a type that a sensible subjectivist position (as described in §2) cannot deliver. In order to make this claim, I start by distinguishing two senses of normative necessity (in §4): *Evaluative normative*

*Thanks to Derek Baker, Catharine Diehl, Barry Maguire, Ville Paukkonen, Karl Schafer, Lucas Thorpe, Teemu Toppinen, Ken Westphal, and audiences at Boğaziçi University and the University of Helsinki for useful comments. Thanks to Philippa Foot—at least the timeslice of her prior to her turn to virtue ethics—for providing the inspiration and the title for the piece.

necessity holds that normative reasons are invariant under change in context from our normative standpoint. *Ontic normative necessity* holds that normative reasons are invariant under change in context *and* normative standpoint. Many of the putative costs of normative contingency are grounded in failures or perceived failures of evaluative normative necessity. Losing evaluative normative contingency would be significant cost, but there is no reason to think that a reasonable subjectivist view will be evaluatively contingent.

Moreover, we do not need to interpret claims of normative contingency as expressing evaluative normative contingency. Claims of ontic normative contingency are perfectly clear claims as to the existence of reasons from within a standpoint. And it is plausibly ontic contingency, not evaluative contingency that subjectivism requires. In response to this general thought, some have argued that there is no sensible distinction like that between ontic and evaluative reasons ((Blackburn 1998; Dworkin 1996)). I think this is a mistake; I define below the notions of evaluative and ontic necessity in terms that make sense on any subjectivist view. If, in spite of this, we cannot distinguish the two, we are owed an argument as to why this is. On the sort of unabashedly subjectivist views I am interested in defending, it seems beyond clear that they can be expressed.

Since subjectivist views are ontically contingent, I also need to show that the costs of accepting ontic normative contingency are not as high as might be supposed. Recognizing that our reasons are contingent, in the ontic sense, does not undermine their normative significance for us or for evaluating others (argued for in §5), we can perfectly well make sense of why we engage in normative argument—at least in the majority of cases—even given ontic contingency (§6), the role of normative judgments of hypothetical cases for our actual normative planning makes even more sense with evaluative contingency than it does for ontic contingency (§6), and, finally, we can assuage the worry that certain intuitively false counternormative conditionals come out true on subjectivist views (§8). Importantly, in a large majority of cases where the agents of a considered context are relevantly similar to ourselves, the distinction between ontic and evaluative reasons more or less collapses; this fact partially explains why we have been so quick to assume that ontic normative contingency is problematic. If ontic normative contingency were true when we consider cases of agents very similar to ourselves, then it would again be nearly the end of the world. But, very plausibly, it is not.

This essay is meant to be a palliative, not a cure, for the costs of normative subjectivism. It cannot be said that there is no way in which a subjectivist view conflicts with some of our intuitions about normativity; clearly it does. But we can alleviate the symptoms if not cure the disease. Much of the palliative work is underwritten by what I think of as one of the central insights of (Foot 1972). This is the insight that the attitudes, desires, conventions, and endorsements that typically make up our normative outlook are *stable* in the sense that it would take a quite radical departure from how we are psychologically, culturally, etc. in order for our normative outlook to significantly change. I will not argue for this directly—that would be sociological project, not a philosophical one (though see §6-7 below and, of course, (Foot 1972))—but there are good reasons to believe it is true. And, if it is, then most apparent change in our moral outlook by agents we view as like us is *derivative* in the sense that it is some local non-normative fact that is disagreed upon or which differs from context to context.

For example, burying your dead and eating them could—and in the famous example, plausibly are—different ways of showing respect to the dead; the parties to the dispute need not be viewed as disagreeing about whether or not we ought respect the dead. The stability of our normative beliefs couples neatly with the pervasive, but widespread, mistake of thinking that our norms must be universal in some sense (see below). This pair helps to explain *why* we have tended to think that our normative reasons must be necessary, universal, or what have you. Given that our normative beliefs are as stable as they are and that we expect them to be universal, it is to be expected that we find it difficult to conceptualize what it would be like to have a radically different set of normative beliefs. But, if this is correct, then we should not be overly concerned about, say, intuitive reactions to certain bizarre normative situations as they are not clearly probative. If we feel that an agent is acting irrationally or counternormatively in such situations, this may just be a projection of our normative beliefs and reactions onto the agent in that situation—which begs the question against the subjectivist about normativity. What we need is a reaction from *inside* the normative outlook of the agent in that circumstance. But that outlook may be nigh impossible for us to conceptualize. And even if we could get such a reaction, this wouldn't show that we should take our evaluation of their irrationality to bear on our own normative planning.

This strategy is broadly similar to an older—and I think unjustly neglected—response to the rule-following problem. This strategy, due to Barry Stroud, picks up on Wittgenstein's example of the Wood Sellers: people who measure the

amount of wood by the area of the ground covered by it. Stroud argues, convincingly in my view, that the only way to actually conceptualize such an example involves importing to them, unfairly, many of our own norms about how to measure and count. But evaluating such a case under such a presumption does not show that their practices are irrational—it just shows that it is irrational to measure wood that way *while* accepting our usual norms about how to measure and count. It also echoes a useful discussion of similar issues by Sharon Street (Street 2009, Lessons 4-5, 10-11) that we need, when considering hypothetical cases, to spell them out completely and not assume that we can make some change in their normative outlook without it spreading throughout the rest of our normative standpoint.

So, if Stroud and Street are right, then the idea that we can conceptualize what it is like to obey this strange convention from the *inside*, so to speak, is fraught with difficulty. But, likewise, we might think that the idea that we can conceptualize what it is like to be governed by some very strange set of desires—say, Future Tuesday Indifference (Parfit 2001) or an intense desire to spend all of our time counting blades of grass—is likewise fraught with difficulty. We'll return to this below. It is important to note that I do not mean to give a knockdown defense of subjectivism here. My purpose is far more minimal—I want to remove a *prima facie* stumbling block to developing subjectivist theories that derives from the sense that our normative judgments must be necessary. I agree, here and below, that if we had evaluative normative contingency, then that would be severe cost to the view. But, as will emerge, it is entirely unclear what the costs of ontic normative contingency are. The goal then is to strengthen the case for subjectivist views that can be given (and has been given, for example, in (Sobel 2009), (Harman 1975), (Dreier 1990), and (Velleman 2013)) by showing that a putative cost of the view is mainly illusory. It is not to address other arguments against subjectivist views that could be given; that is a matter for another occasion and, somewhat, for other authors. Finally, note that many of the points and techniques I bring to bear below are well-known. My aim is not to give an entirely new defense of subjectivism, but rather to draw together a number of threads with the aim of showing how they can be put together in such a way as to undermine one standard and looming objection to subjectivist accounts.

2. Normative Subjectivism

The particular subjectivist views I have in mind share the following characteristics. First, they are characterized by an explanatory biconditional like the following:

We have to reason to ϕ if and only if (and because) ψ

where ψ is some condition or state of affairs that is essentially dependent on *contingent* features of me or my moral community.¹ Given a view characterized by such a biconditional, our reasons are contingent on our circumstances in a way that many have found objectionable. Such views, at least in the versions I will consider, also allow:

It is correct *in some sense* to say, of a , that they have reason to ϕ if and only if and (because) ψ

where again ψ is a condition or state of affairs that is essentially dependent on contingent features of me or my moral community. Note that this means that it is possible—though not required—that I, a member of moral community A, can correctly say of a member of moral community $B \neq A$, that B has reason to do ϕ *even if* they cannot say this of themselves. I can, therefore, correctly condemn the moral practices of another community even if their practice is not correctly condemnable from within. Not all subjectivist views hold this, but ones that deny it have a much harder time satisfying the demand for universality that normative judgment seems committed to. I discuss this evaluative notion in detail below.

The following four sample accounts of reasons (for me to avoid harming innocents, say) are paradigmatically subjectivist:

- (H) I have reason to avoid harming innocents just in case and because I want/desire to do something which not harming innocents is identical to or brings about.
- (R) I have reason to avoid harming innocents just in case and because not harming is part of or a consequence of a system of norms I subscribe to.
- (E) I have reason to avoid harming innocents just in case and because I disapprove of harming innocents.²

¹Note that this means that Kantian ethics does not satisfy this biconditional—for Kant, our reasons essentially depend on facts about us (qua rational agent), but such facts are not contingent.

²I am taking (E)—obviously intended to be an expressivist view—as subjectivist. This is the

(J) I have reason to avoid harming innocents just in case and because I do not want to be open for blame for doing something wrong—such as harming innocents.³

Note, importantly, that these are not claims about the *meaning* of reasons claims, but rather claims about *what it is* for there to be a normative reason for me to do something.⁴ And in each of the four, the explanans is contingent in at least one of the two senses I will sketch shortly.⁵

Since the normative facts on these views depend on contingent facts about us, normative facts themselves are plausibly contingent. At least, on the most natural way of spelling out the above views, suitable shifts in our attitudes, norms, or desires will generate a shift in the normative facts. Note, again, this is not the claim that, on these views, we need to endorse claims like “If I were to want to kick dogs, kicking dogs would be okay”. We need not (see §8). Rather, I am making the *metanormative* claim that, if these views are correct, then if we were to desire to, approve of, or accept norms permitting kicking dogs, we would have reason to kick dogs. This is a different question from whether we ought to endorse this claim from *within* our perspective.⁶ I am not concerned with chiseling down the right-hand side of the labeled theories above—of course, for the usual reasons, we want to add hedges to avoid worries having to do with false information, rash approvals, non-endorsed first-order conative states, and the like. I will suppose that some such chiseling can be done for the purposes of this paper. For what it’s worth, my own view is that reasons are relatively cheap and most of the worri-

result of treating (E) as view of normative discourse, but to avoid a commitment to *quasi-realism* of the stripe endorsed by Blackburn and Gibbard. The issue of quasi-realism is too difficult to enter into here.

³For the details of this relatively new view, see (Woods forthcoming).

⁴This means, in principle, that it could turn out that what we should say about reasons and what reasons there are might come apart. More on this below.

⁵Street’s Humean Constructivism (Street 2010; Street 2012) would also count as a subjectivist view, at least given the way I’ve taxonomized things. I will not deal with it explicitly in what follows as I find it difficult to understand whether or not a view is constructivist, but I recognize the similarity of it to, say, R and J. What I say here applies to her view as well.

⁶Again, I put to the side here those, like (Blackburn 1998) and (Dworkin 1996), who suggest that there is only one way to make sense of such claims. Even if they’re right about the views they favor, which I strongly doubt, we can still make this distinction on any self-described subjectivist view. Since my purpose here is to defend subjectivism from contingency and arbitrariness objections, I will thus assume that the above views permit the distinction I’m here articulating.

some cases need swallowing more than they need avoiding.⁷ However, it is too hasty to assume that we will always have sufficient reason to be moral or obey some normative standard, in contrast to (Schroeder 2007, ch. 6). Even though I find the idea that we always have some weak reason to be moral fairly plausible, at least for creatures relatively like us, we should not rule out ahead of time that there could be normative standpoints where there is little if any overlap in what reasons we have. Or, even if there is overlap, that nevertheless there is no overlap in the strength of our reasons to conform to normative standpoints.⁸

I talk here in terms of reasons, not obligations, for no particular reason. I could run the discussion either way, but reasons talk is increasingly pervasive and we can plausibly recapture at least some sense of obligation in terms of what we have most reason to do. Anyways, the contours of an actual subjectivist view are going to be tremendously complicated, but that does not affect the general structural points I want to make. I will henceforth concentrate on view R above so as to simplify the discussion. Nothing turns on this choice; my argument could be made either way. But it simplifies the discussion to concentrate on one view and, for what it's worth, conventional views seem to me to have the most plausibility of the subjectivist options. I presume that conventional facts tell us not only what reasons we have but also what relative strength they have. This is a fairly substantial assumption, but, again, it will do for our purposes here.

3. Normative Necessity – First Pass

That normative facts are simply not arbitrary or contingent is a familiar claim:

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

⁷So, I agree with (Schroeder 2007) that the having a reason to eat my car since I want to have more iron in my diet is a pragmatic phenomenon—such reasons are so weak that it is bizarre to mention them when articulating our reasons to someone else.

⁸For the most part, I will ignore the complication of strength of reasons below. There will often be independent reasons to obey any particular chunk of our normative standpoint (Schroeder 2007, Woods forthcoming). But such reasons are not really what is of concern here, so I will henceforth ignore this fact for the sake of expository simplicity.

with absolute necessity... (Kant 2002, 24)⁹

and, perhaps in addition, that if we recognized normative facts as arbitrary or contingent, then we would lose our reasons to take them seriously as reasons for action:

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

My aim here is to overturn subtle versions of these intuitions, but I need to start with the bluntest of them. Kant, quoted above, seems to be articulating a brute intuition about normative facts. They are *categorical* and *binding* on all rational individuals—at least the normative facts which concern morality. But this brute intuition is hard to defend. As Philippa Foot pointed out, the categorical form of (moral) normative judgments is insufficient to validate this intuition. Facts in the obviously contingent etiquette sphere likewise have categorical form. We do not say “take your feet off the table if you don’t want to be rude” or “you ought take your hat off unless you want to be impolite”, but merely “take your feet off the table” and “you should take off your hat”.¹⁰ But what then justifies our intuition that normative judgments are necessary?

It cannot be the thought that we do not find diversity among normative opinions. Diversity in moral opinion, at least about large swathes of moral opinion, is simply a fact of life. And 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 we can best care for ourselves. Even diversity in epistemic opinion is rife: consider the endless Facebook debates about the evidential status of testimony.¹¹ So what could underly the supposed necessity of normativity? At least at first glance, it seems that nothing does. And, perhaps, that our desire for normativity to be necessary is an articulation of our fear that if it were not, then our judgments of others and use of normative judgments as a guide to behavior would be simply a mistake. Smith’s argument above seems to be an articulation of this fear. Philippa Foot also articulates the potential motivation here well:

⁹Lightly modified by Catharine Diehl.

¹⁰See (Foot 1972). Others, such as Miss Manners, have taken this to imply that etiquette facts simply are moral facts (Martin and Stent 1990), but this is at serious variance with our intuitive judgments about etiquette facts where they seem to have only local force. They seem to be paradigmatic instances of *merely formal* oughts.

¹¹*Nota Bene*: such debate is not typically carried out in philosopher’s jargon.

But this is not what is urged by those who think they can close the matter by an emphatic use of “ought.” My argument is that they are relying on an illusion, as if trying to give the moral “ought” a magic force.

This conclusion 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, 315)

But why should we take this fear to be indicative of something actually to worry about? We know full well that aesthetic taste, epistemic standards, norms of etiquette, and the like are relative. Or, anyways, they are widely believed to be relative, at least when reflecting on them, and that is enough for the point to be made here.¹² Reflecting on this does not undermine their normative significance *for us*; far from it.¹³

And it does not show that it is not reasonable or even demanded that we entirely purge ourselves of our habit of engaging in normative evaluation of others from within our own perspective. As I will suggest below, this form of evaluative necessity seems to play a useful pragmatic role in regulating the normative outlook of others and in stabilizing our own normative outlook against various changes. Even though I know that it would seem fine to me tomorrow to count blades of grass for the next 8 hours (were I to take a few tabs of acid and read the unjustly neglected essay “On the Joys of Counting Blades of Grass”), my current repugnance at the thought of doing so serves to keep me on the straight and nar-

¹²For what it’s worth, it seems to me common knowledge that aesthetics and etiquette are subjective; arguments that they are not tend to privilege notions of artistic skill and usefulness, respectively, that are better thought of as independent justifications for having the conventions or aesthetic standpoint that we in fact have (see REDACTED for relevant discussion of this sort of confusion.) Anyways, the issue is too complicated to enter into here. I will just presume that aesthetics and, even more so, etiquette are subjective in the relevant sense.

¹³I here disagree with James Lenman’s otherwise admirable (Lenman 1999). Lenman suggests “If a side-effect of this medicine is that I lose my taste for Coca-Cola, I could not care less”. But, typically, we do care about such things; our current tastes inform our evaluation of potential future tastes. I do wholeheartedly agree with his claim to moral generality: “However, such 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.”

row.¹⁴ Of course, such a role can misfire (I still wince at the thought of certain vegetables even though I know that with a minimal effort I could get myself to like them and it would be good for me to do so), but so what? Diachronic stability of our normative outlook, even if it goes astray now and then, is nevertheless typically a good-making feature of normative outlooks.

Smith's claim about undermining, taken as a descriptive claim about our normative outlooks as applied to us, simply seems false. And as a normative claim, it seems unmotivated—it is entirely unclear why recognition of the fact that my normative reasons are sourced in my moral standpoint (and that that standpoint is somewhat arbitrary) *should* make me less apt to apply them to myself. After all, recognition of the fact that my aesthetic or etiquette reasons are sourced in my aesthetic or etiquette standpoint (and that such standpoints are somewhat arbitrary) makes doesn't undermine my reasons to apply these standards to myself. So, if we understand Smith's objection in terms of undermining of the normative significance of our normative reasons for ourselves, it seems mistaken.

But, plausible as Foot's remark is, perhaps there is a more generous thought in the vicinity. Perhaps the relevant thought is that normative facts are reason-providing for not only ourselves, but others, regardless of whether they share or do not share our normative beliefs and that this is part and parcel of why we engage in normative theorizing to begin with. That is, perhaps we should understand Smith's worry in terms of undermining the normative significance of our moral reasons for others. Following this thought out, perhaps it is part of the functional role of normative theorizing that we be able to criticize and correct the actions of others. Lenman, discussing Smith, gives a gloss of what this might mean:

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

This sort of argument from the usefulness of disagreement seems to me to have some force. It would 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 arbitrary others. We do

¹⁴Credit where credit is due—I first read the counting blades of grass example in a paper by David Sobel.

not take etiquette facts to have such a basis—my judgment that it is rude to expose my feet gives me no reason to fail to take off my shoes in a Turkish household; their sense of the rudeness of failing to respect teachers 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 must or, at least should, change accordingly, at least when we think carefully about it. But we are not so aware in the case of morality and this may be due to differences in the role of moral, normative, and etiquette theorizing.

So perhaps this is the essential bit of the necessity of normativity; 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. But the relationship of this putative fact about the necessity of our conception of normativity to the facts about our normative reasons themselves needs to be further spelled out. There are many ways of satisfying the above facts about normative theorizing and not all of them require the *metaphysical* or *conceptual* necessity of normativity—something far weaker will suffice. The fact of the matter is that we still very much take seriously the idea of disapproving of someone for failing to share various outlooks *we full well know* to be less than fully objective. I hate stirrup pants, think vegetarian Chili is a travesty, and disapprove of avocado. And I think less of people who do not share these tastes. But I know that these are more or less the arbitrary by-product of arbitrary whim; and though I don't think they are fully arbitrary (I can give reasons for my tastes), I nevertheless recognize that I might very well have had different ones. But why think that this should make me more tolerant? Tolerance of despicable aesthetic tastes I very well might have held is no part of my aesthetic outlook, just as tolerance of despicable normative outlooks I very well might have held is no part of my normative outlook. We will turn now to some details of this sort of response.

In this section, I have discussed two straightforward objections to the idea that normativity could be subjective. First, the Kantian claim that normativity is universal and categorical and second the Smithian claim that if we were to come to believe that our normative reasons were subjective, we would accord them no normative significance for us. I argued that the Kantian claim was plausibly undermined by Foot's discussion of the categorical form of etiquette judgments and that Smith's claim seems to jar with the fact that many types of reasons, such

as aesthetic reasons, were known to be subjective and still did not lose their action-guiding significance for us. I then suggested that the better form of this objection had to do with undermining the role of our normative judgments in criticizing others, discussion of which we will now turn to.

4. Two Notions of Normative Necessity

Let's distinguish two ways of understanding the necessity of normativity. First, and most clearly, we may view normative reasons as holding for each other when *evaluated at our own perspective*. That is, when I assess what you have reason to do, I may draw on my normative beliefs. So when you suggest that eating babies for a mid-afternoon snack is permissible, I may draw on my quite reasonable belief that we ought not to harm innocents in satisfying mild hunger cravings and present this to you as a reason for you to avoid munching down on wee Davey. Not only this, but if you reject this reason, this gives me no reason to revise my opinion that you ought not mow down on the child and no reason to revise my opinion that you are doing something you have significant reason not to do. After all, your normative beliefs are mistaken.¹⁵ Call this the *evaluative* sense of normative necessity: a claim that we have reason to do something is evaluatively normatively necessary when it holds no matter what the non-normative facts may be *when evaluated from our normative perspective*. Consider, for example, a normative perspective that assigns significant weight to utility maximization. For such a view, "we have reason to do the thing which is utility maximizing" will be evaluatively normatively necessary, whatever the views of the agents in the context we evaluate. If their normative viewpoint tells them to do something which does not maximize utility, they're simply wrong (at least by our lights.)

Second, and slightly more obscurely, we may view normative reasons as holding for each other when *evaluated even within each others perspectives*. That is, we may view normative reasons for me, in context α , and for you, in context β , as holding regardless of whether we evaluate them from within α or β . We thus may view normative reasons, when correct, as holding equally for each other within our relevant perspectives. Given this way of taking normative reasons, when you reject my quite reasonable belief that we ought not to harm innocents in satisfying

¹⁵(Schafer 2014a) calls a subjective view cast along these lines *perspectivalism*. Such a view promises to vindicate the necessity intuition, cashed out in terms of extensional adequacy of our normative judgments across all sorts of seemingly deviant cases, without abandoning the intuitive thought that our normative reasons are sourced in our normative perspective.

mild hunger cravings, then insofar as I think you are speaking sincerely, carefully, and with a reasonable grasp of the non-normative facts concerning the case, I ought to accord your rejection *prima facie* weight in my deliberations about what you ought to do *vis a vis* your standpoint.¹⁶ Call this the *ontic* sense of normative necessity: a claim that we have reason to do something is ontically normatively necessary when it holds no matter what the non-normative facts may be *when evaluated by the normative perspective embodied in those non-normative facts*.¹⁷

To fix ideas, let's use some machinery from (Einheuser 2006). Let a "world" be a pair $\langle c, n \rangle$ of a *context* 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. I take a context here to be 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, or what have you.¹⁸ We will say that a world *gives rise* to a reason for an agent a to ϕ when a system of norms, applied to the non-normative descriptive facts, says that we have reason to ϕ . We'll write this as $\langle c, n \rangle \models R(a, \phi)$.¹⁹ Following Einheuser, we say that a system of norms is *grounded* by a context (writing this as n_c) when the conventional practices that find their expression in behaviors and psychologies described in c are n . So, for example, the actual context $c_{@}$ grounds a system of norms n where we have decisive reasons not to kick dogs (i.e. in the conventional practices embodied in $n_{c_{@}}$, we have decisive reason not to kick dogs.)

Now we can vary our actual context $c_{@}$ and our actual norms $n_{c_{@}}$ in two ways, obtaining two notions of necessity. What I've called *evaluative* necessity can be

¹⁶This is not to say that we do this often; such disputes are relatively uncommon instances of "normative" deliberation. See §7 below.

¹⁷This way of distinguishing two notions of reasons is similar to the two-dimensional interpretation of expressivism given by (Peacocke 2003). Peacocke uses this distinction to attack the mind-independence of Blackburn-style quasi-realism. This attempt is interesting, but outside the scope of the present essay. For some related discussion, see the below section on counternormative conditionals.

¹⁸We are using fragments of possible worlds instead of just possible worlds themselves to allow that different communities in the same world may have different normative behaviors. If this jars, make the simplifying assumption that there is only one set of behaviors at any world and let a context be a possible world. For the point I'm making here, nothing much matters on spelling out the tedious detail.

¹⁹We will simplify in the relevant cases 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.

defined for the conventional account as follows

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

That is, it is evaluatively necessary, at a world $\langle c, n \rangle$, that a has reason to ϕ in a situation γ just in case that a has reason to ϕ in situation γ at a world $\langle c', n \rangle$ which differs from ours only in context (i.e. non-normative facts), but where a and γ exists at context c .²⁰ That is, when evaluating by the conventional norms of our context, no matter what context we look at, a has reason to ϕ in γ . Now, clearly, very few, if any, particular claims like this will be true since most of our reasons are partially dependent on circumstance. But we can find a plausibly evaluatively necessary claims easily enough; consider “Charlie has no reason to murder someone just for fun” or “no one ever has a reason to be cruel merely for fun”. Evaluated by a stringent version of our own moral perspective, 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 evaluative contingency. Our norms may not express the idea that others have reason to do what their conventions suggest, we may disapprove of others following their conventions, etc.²¹

Defining *ontic* normative necessity is also straightforward:

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

That is, it is ontically necessary, at a world $\langle c, n \rangle$, that a has reason to ϕ in γ just in case that a has reason to ϕ in γ in any world composed of a context c' and the norms which are grounded by c' . That is, if a has reason to ϕ in any non-normative context where their actions are evaluated by the conventional norms of that context. This and the above definition of evaluative necessity can be extended to cover all complex claims involving R in the obvious way. Subjectivism of the above stripe does entail ontic contingency so long as different contexts embody different systems of norms, psychologies, or what have you.

We can expand on Einheuser’s terminology to define two sense of reasons. Let us say that a , in their context c_a , has an *ontic* reason to ϕ in γ just in case

²⁰I will take the restriction that a and γ exists at c' as understood henceforth.

²¹Note that in some cases, there will be derivative reasons to do such and so that derive from an explicit reference to the local conventions—when in Rome and all that. But the base facts for a subjectivist theory *need not* make any such reference.

$\langle c_a, n_{c_a} \rangle \models R(a, \gamma, \phi)$. And, fixing a set of norms n , we will say that a has an evaluative _{n} reason to ϕ in γ just in case $\langle c, n \rangle \models R_e(a, d, \phi)$. 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 applying n to c gives rise to a reason to ϕ in γ . Given these definitions, it's easy to see that a 's reason to ϕ in c is ontically necessary 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 necessary just in case n yields a reason for a to ϕ in γ given any context c' .²²

Now, as mentioned above, many of our reasons will be derivative. Consider, for example, our reasons to not kick dogs. Presumably these reasons are not fundamental, being a consequence of a norm inveighing against causing pain in the context where dogs feel pain when you kick them. As such, they are not necessary in either of our defined senses. But, presumably, the claim that we have reason to not cause pain is evaluatively, even if not ontically, necessary.²³ At least on many plausible articulations of the views mooted above, we will have fundamental claims about our reasons which will be evaluatively necessary. And, given these views, we can articulate a non-trivial notion of evaluative supervenience. We will write $\gamma \cong \beta$ when two situations are alike in all non-normative respects.²⁴ Then, we can define evaluative supervenience as follows:

$$\gamma \cong \beta \Rightarrow \langle c, n \rangle \models R(a, \gamma, \phi) \text{ iff } \langle d, n \rangle \models R(b, \beta, \phi)$$

I am assuming, for the purposes of this definition, that the relevant non-normative properties of a and b and the relevant non-normative features of γ and β are part of what is going into the situation variable—so, for example, that a is really pain sensitive and b not would destroy the similarity of γ and β as would the fact that a 's ϕ -ing would go unnoticed and b 's not, given other features of their context. Now, note, this is just an example supervenience claim. Tighter versions could be articulated which took into account further normative features of the respective

²²This distinction could be articulated in a number of different ways, but given that I've used Einheuser's machinery in order to distinguish the types of necessity, I will follow suit in distinguishing the types of reasons. Translations back into one's favored notation or machinery should be straightforward.

²³I here assume that pain is unpleasant and to-be-avoided.

²⁴I am here assuming that there is some way to articulate a difference between normative and non-normative features of a situation. This is a difficult problem, to be sure, but one not unique to my projects. So I will bracket difficulties here and assume that some such distinction can be drawn.

situations. For example, perhaps for some systems of norms, there's a "buy 5, get 1 free" clause which excuses us from reasons to continually do good things. Or, perhaps, we have additional reason to help those who've been wronged before. But it's possible to modify the definitions accordingly; for my purpose here, it's sufficient to give an example. The above definition can be extended in the obvious way to cover complex claims involving R .

The example is not entirely arbitrary. It can easily be explained on the basis of features of our normative standpoint. If, as we've assumed above, our normative standpoint consists of fundamental principle like claims about our reasons and if derivative reasons follow from these norms given a particular situation, then if there is no change in the relevant features of the particular situation, there will be no change in the relevant reasons that follow from our fundamental norms. For particularistic accounts, like some construals of disapproval-centric or desire-centric accounts will be, it is slightly more difficult. Still, it is plausible that there is an explanation of why we have the reactions we do in certain cases, and this explanation is likely to draw on non-normative features of the relevant situation and non-normative features of our normative standpoint. Given this, it is extremely plausible that some version of a supervenience account will hold. Finally, as noted by Blackburn early on, expressivist and relativist views can embed, on top of their situation-targeted norms, general coherence norms like "treat like cases alike". Given these, we will have a rather direct explanation of why supervenience claims, in some form or other, will hold.²⁵ Finally, note that although we can define ontic supervenience easily enough

$$\gamma \cong \beta \Rightarrow \langle c, n_c \rangle \models R(a, \gamma, \phi) \text{ iff } \langle d, n_d \rangle \models R(b, \beta, \phi)$$

it is entirely mysterious why it would hold on a subjectivist account. It seems inevitable that it will fail, given what it is to be a subjectivist account in the first place.²⁶ Difficulties explaining ontic supervenience are notorious,²⁷ but the issue is too large to enter into here. It suffices to show that there's a non-trivial form of supervenience available to the subjectivist which is easily and naturally explicable.

²⁵The relevant strength being determined by the account of what it is for cases to be alike. This matter is too extended for treatment here; I hope to return to it elsewhere.

²⁶Of course, this is only presuming we do not include the grounding features of the context in the situation variable. We will ignore this possibility since doing so would trivialize the supervenience property.

²⁷See (Väyrynen forthcoming) for a useful survey of such arguments and (Blackburn 1988) for the classic case against normative nonnaturalists.

5. Subjectivism and Criticizing Others

Both types of reasons and the articulations of evaluative and ontic necessity and supervenience make sense on their own terms. After all, I've just defined them in a way which is perfectly coherent, given a conventionalist setup, and it's easy to see how to adapt this set-up to any of the alternatives mooted above. But the explanation of why we hew to each form of necessity, if indeed we do, differs quite dramatically. To see this, consider again the case of aesthetic judgment. All the above definitions go through for the aesthetic case, though the details of our aesthetic standpoint will differ somewhat from that of our normative standpoint. Importantly, it is *known* that tastes are relatively subjective and that people, even in small communities, can differ quite wildly about their tastes. In fact, this is known rather dramatically to me as I am no fan of chocolate; I find chocolate bearable, but generally unremarkable. This is rather strange to most people. And assessed from their gustatory standpoint, I do have reasons to eat more chocolate and, perhaps, even to develop a taste for it. "Assessed from their gustatory standpoint" here means evaluating what gustatory reasons I have according to their tastes—i.e. looking, from their standpoint, at what evaluative reasons I have—and the output of their assessment of my reasons to eat chocolate will diverge from my own predilections. Correspondingly, the evaluative reasons I have, from their standpoint, will differ from the evaluative reasons I have according to my own gustatory standpoint—i.e. the ontic reasons I have. In other words, it is plausible that I have an evaluative reason to eat chocolate according to the gustatory standpoint of the average individual, but no ontic reason at all.

Consider likewise the joke American's make about British folks driving on the wrong side of the road. This is a joke—kind of. The funny thing about the joke is that we *really do* feel like British folk drive on the wrong side of the road, even though we know quite well that it is entirely arbitrary which side of the road to drive on. The intuitive force of the approval of left-driving over right-driving remains, even when we know quite well that the choice of side is arbitrary. Even in cases where we know quite well that how we solve a particular coordination problem is more or less arbitrary, there is a strong temptation to project our solution onto others and view badly their solutions.

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. In other cases where we

have explicitly subjective reactions, such as cases of color perception, we know that recognition of this undermines claims to correctness and incorrectness.

The reason this ambivalence is harmless is that once we bring other perceptual systems into view, the provided they are equally discriminatory, we lose any very robust attachment to the idea that ours is right and theirs is wrong. Similarly we do not maintain sceptical fears that perhaps our sense of smell, or sense of colour, may in general be letting us down, so that perhaps things really smell differently from the way we smell them, or have different hues from those we see them as having. 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, contra Blackburn, that many of us feel a drive to criticize others aesthetic sensibility *even though* we know full well that tastes are subjective—or, anyways, my experiences being a chocolate deviant suggest quite strongly that folks do so. 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 tend to be plastic and, from your point of view, chocolate really is delicious. Continued experience with a particular thing, be it chocolate, wine, alcohol, or the like, tends to breed some taste for it, from your perspective chocolate is high on the valuable tasting things, and, of course, some convergence in our tastes is a useful and desirable thing. And, perhaps, we would 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 sort of explanation cannot be said to hold for ontic aesthetic reasons. It is 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. Clearly some of us do not. And given that I have no desire for chocolate, there is no gustatory reason for me to pursue the modification of my desires that continued exposure to chocolate might provide. And given that I am assessing my gustatory reasons from within my gustatory standpoint, if I have any reason to develop an affection for chocolate, it is the prudential reason of not standing out, not the gustatory one of developing a finer aesthetic sensibility. Yet this type of aesthetic irrationality is a pervasive feature of our 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. The *explanation* of why we might hold onto ontic aesthetic necessity thus seems to not justify holding onto it at all, unlike the case of evaluative aesthetic necessity—it seems rather a case of projecting some kind of mistaken universality on a bit of our experience we know, on reflection, to be non-universal.

As with aesthetic reasons, so too with normative reasons. We can make sense of evaluative normative reasons and thereby evaluative normative necessity. And the reasons given above for evaluating the aesthetic reasons of others hold even more 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 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 others actions, even when they clearly don't share our fundamental practical norms.

Applying *our* normative notions to *your* actions yields a definite verdict and that definite verdict makes sense—even if it is not obviously the verdict you should share.²⁸ Of course, again, our normative notions might be tolerant of differing normative opinions, but as I said above, sensible subjectivist views can and should accept that we can evaluate someone else's actions independently of their normative views. At least if they are to be at all plausible as a descriptive account of our actual normative practice. It's here that we find the truth in Smith's idea that viewing our own perspectives as arbitrary would undermine them (Smith 1999). If we viewed our reasons as holding *evaluatively* only for us or those with extraordinarily similar normative practices, then we would 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. And given that exertion of social pressure by means of explicit normative pronouncement is reasonably effective as a tool of social coordination, we can see good reason to engage in this practice of evaluating others.

But, again like the aesthetic case, it seems simply to be a mistake to think that we have reason to accept ontic normative necessity, at least if we refrain from

²⁸Consider the similar discussion of defined notions of evaluation in (Schafer 2014a).

begging the question against the subjectivist. Why should we expect, absent a commitment to some form of normative realism, that reasons from within our normative standpoint will hold with equal generality within the normative standpoint of others, except by happenstance? Ontic normative necessity is simply implausible in the context of a subjectivist theory like those outlined above and it seems dialectically unreasonable to hoist such a conception on the subjectivist. Evaluative normative necessity, plausibly. Ontic normative necessity, no. And this seems to track ordinary responses to moral relativism. Folks get upset—quite reasonably so!—at the thought that they cannot judge the actions of others when faced with suitably horrible examples; they worry far less if at all about whether or not there are *really* reasons from within the perspectives of others to do horrible things.

In this section, I have argued that recognition that our reasons are sourced in our normative standpoints and that standpoint can differ does not undermine our tendency to bring our reasons to bear in criticizing others. In particular, I have distinguished two sense of normative reasons—evaluative normative reasons and ontic normative reasons—and defined two notions of normative necessity—evaluative normative necessity and ontic normative necessity. I then argued that what we need in order to make sense of our bringing to bear our normative reasons on others is evaluative normative necessity. But evaluative normative necessity is completely consistent with subjectivism about normative reasons in the sense we are concerned with. This leaves open the question of how and why subjectivists should engage with others about what they have reason to do; we turn to this now.

6. How Subjectivists Should Disagree

The above section argues that we are often willing to evaluate others normative standpoints by means of our own even when we know that those we evaluate do not share our point of view. I've also suggested that it is reasonable to do so (and more in this vein can be said: see §7). But, even given this, there remains the question of how to understand another descriptive fact about our normative practice; we do not only judge others from afar, we also frequently engage in explicit argument with others about normative matters and attempt to convince them to modify their normative standpoints or adopt reasons claims that we accept. But, if their normative standpoint is significantly different from ours, then this practice might seem to make little sense. *De gustibus non est disputandum*, after all, it might seem. But, if what I've said above is plausible, then we should not take

that Latin phrase that seriously. We dispute about taste all the time.²⁹ But, if we believe that our normative standpoints, or even our aesthetic standpoints, are not uniform, then why would we do such a thing?

In answering this question, it's useful to set aside at the outset a class of cases where it really is not useful to argue about normative reasons. If we know, in advance, that someone has a coherent and deep commitment to a normative standpoint drastically different than our own, one which has little to no overlap with ours, then it really does not seem very useful to dispute with them. This point echoes a neglected discussion of moral argumentation in (Ayer 1946). Ayer describes the point of moral argumentation as proposing various “deeper” claims until our interlocutor latches onto one of them. Punching him is okay? Well, do you think hurting is fine? No? Do you think causing pain unnecessarily is fine? No? and so on. Ayer goes on to suggest that if this procedure does not work, we tend to abandon the argument. And, presumably, if we can see in advance it will not work—our potential interlocutor is so obviously different from ourselves—then it seems unreasonable to start arguing with them at all. Of course, Ayer is neglecting the role of argumentation in *irrational* change of belief—we can bring people to agreement with our views by means of mere argumentative peer pressure, sometimes, and it is not obviously unreasonable to do so (see also (Street 2009, lesson 10) for an articulation of this point.) But, when someone has suitably robust views and a suitably argumentative nature, then it really does seem slightly pointless to argue with them. Kate Manne puts this point well

Suppose, once more, that I am trying to get someone to alter his behavior. For example, to borrow Williams' well-known example, 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. He doesn't care about improving his wife's well-being, his marriage, or even his own lot on this score (assuming, as is plausible, that the

²⁹I've always preferred the punny rendition of the phrase in English anyways. There's no *ac-counting* for taste strikes me as a nice way of expressing that one searches in vain for the reason someone likes lettuce or Radiohead.

well-being of the two partners is not wholly unconnected). He is not interested in being a good husband in the abstract. Nor does he care about acting more kindly and considerately in general. He is not even moved by the thought that his wife is a fellow human being, who will be hurt by the shabby way in which he continues to treat her. And so on and so forth. We are not getting anywhere.

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...Entering into this exchange, I had hoped to have what we sometimes call a 'rational conversation.' (Manne 2014)

When the desire for a rational conversation is frustrated, it seems pointless to do anything but try to browbeat one's interlocutor into moral compliance.³⁰ Likewise with the perfectly coherent teleo-Catholic or any others of a similar stripe. It's pointless to try to argue with them as their and our normative standpoints differ, at least in the relevant cases, to simply too large a degree. But put such aberrant cases to the side for a moment. If we think that our interlocutor is reasonable—if we think, to borrow Manne's expression, that they are open to a 'rational' discussion—then I think we assume that they and we agree, to a large degree, about what sort of considerations matter for making claims about reasons. In fact, I think it is plausible that such overlap is required to make sense of rational discussion at all.

(Carballo and Santorio forthcoming) argue that in order for rational communication about normative matters to have a point, both participants must take as part of the presupposed common ground (in the Stalnakerian sense (Stalnaker 2013)) that there are norms—either derivative or fundamental—that are not ruled out in advance of each other and that participants in normative communication hold that they ought to come to some convergence.³¹ The Stalnakerian picture typically is

³⁰Manne is concerned with cases where motivational internalism fails, but we can see the point more generally. It's often pointless to try to get someone to change their mind.

³¹They go on to argue that both participants should agree that there is a unique convergence point, but this seems overly strong for many of the subjectivist views mooted here. If this were true, however, it would just make my point all the stronger. Their general argument is pitched towards understanding how expressivist views can understand normative communication, but the

understood in terms of a somewhat metaphorical notion of self-locating belief, but it is better to understand normative communication in terms of *coordination*. In particular, we should understand communication about normative matters in terms of coming to agree on a norms or attitudes we can share and usefully apply in planning our actions and evaluating the world around us—in particular, both participants in a discussion sharing the view that there is significant reason to come to agreement and that it is possible, at least to some extent, to do so. If this is right—and I suspect it is—then we should assume that our interlocutor and we share some part of our normative standpoint at least in an idealized conversation where the participants have no false beliefs and have committed no factual or logical errors.

To put the point more precisely, what we must presuppose as common ground between us and our interlocutor is (a) that there is some part of our normative standpoint n (call it c) which overlaps with theirs t and (b) that we both accept that we have reason to bring our normative standpoints in line with each other. If we presuppose this, then we are entitled to also presuppose that

$$\langle c, n \uparrow c \rangle \models R(a, \gamma, \phi) \text{ iff } \langle c, t \uparrow c \rangle \models R(a, \gamma, \phi)$$

That is, restricting our norms to the overlapping sections of our normative standpoints, there is a reason for a to ϕ in γ by my lights if and only if there is a reason for a to ϕ in γ by your lights. But if this is right, then in engaging in discussion 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 means of inviting our interlocutor to add a norm, derivative norm, disapproval state, or what have you to the common ground and make the appropriate modifications to one's overall normative standpoint. Since the shared background is going to be 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 of how normative communication might work. First, suppose that you and I are arguing about abortion and we agree (a) about many of the non-normative details, (b) that pain and suffering is bad, we have reasons to

point is more general.

avoid taking life, etc, and (c) that we have reason to make our view more coherent and natural when possible, but neither of us have a stance on whether or not third-trimester abortions are permissible. I may propose that we have reasons to avoid third-trimester abortions on grounds that fetuses are fairly developed by that point and the likelihood of suffering is non-negligible. Since we agree on a and b, presumably I should come to accept this claim, add it into the common ground, and thereby expand the overlap of our normative standpoints. Second, suppose that the case is as above, but I think we have reasons to avoid third-trimester abortions and you disagree. I propose as before, suggesting that if you accept b, you *should*—say on abductive grounds—accept that we have reasons to avoid third-trimester abortions. Since we agree on b and c, you may come to accept that we have reasons to avoid third-trimester abortions, retracting your earlier view.

Obviously details need to be fleshed out here, but the general picture seems to make sense and is more or less sufficient for the point I want to make. If communication about normative matters looks like what I have just described, then we can understand why we do not need full-scale ontic necessity to make sense of normative communication. We do not even need full scale evaluative necessity for this; rather, all we need is that there we presuppose overlap in our normative standpoints. In typical cases, there will be quite extensive overlap. Of course, presuppositions can be false and it might be that our communication really does make little sense given how little overlap there is in our normative standpoints. But the more atypical the case, the more pointless it actually seems to engage in rational discussion instead of engaging in browbeating, the exertion of social and peer pressure, and the like.

In this section, I have suggested, following Ayer and Manne, that not all normative communication is reasonable. Sometimes, if we and our interlocutors are different enough, normative communication is only useful in order to browbeat, peer pressure, or push our opponents to modify their normative standpoint. Following Carballo and Santorio, I then suggested that in order to make sense of communication about normative matters, we need to presuppose that we and our interlocutors share some fragment of our normative standpoint. And if we do so, then we are in a position to offer reasons to our opponent which are sourced in this overlap. But this falls far short of a requirement of ontic or evaluative necessity.

7. Subjectivism and Hypothetical Planning

We have seen that the argument from disagreement does not support believing in ontic necessity. Can we give non-question-begging argument for thinking that ontic normative necessity holds? I can think of one potential argument—an argument from the functional role of normative judgment—which has some pull. Consider Gibbarian views of the role of normative judgment where the role of normative judgment is something like contingency planning for various circumstances. On this view, the role of normative judgment has to do with structuring our own plans of actions; in particular, on figuring out what we would do in various circumstances that attach to our own lives. Much like reading novels, we consider hypothetical and counterfactual situations to test and hone in our actual reasons.

But, given this, we can see why we might want ontic necessity—we want to know that, were we to bring our garnered knowledge to bear on analogous circumstances, the reasons on which we would then act would be genuine reasons for action then as well. And we want our testing of general reasons claims—like that it's wrong to cause severe pain for fun—to be probative for our reasons to act *now*. If the normative reasons are not at least semi-universal, then it is hard to see how our evaluations of what reasons there are to act in certain counterfactual situations could do the relevant job of informing our actual plans of action.

Suppose I know full well that in some other circumstance, my coffee will explode. In that circumstance, I presumably have decisive reason to back quickly away from the coffee (and, perhaps, warn others of the danger). But if I know that this is not my circumstance, it seems irrelevant to my current actions, other than as a philosophical curiosity, that in some other world my coffee will explode. This worry seems less than fully noticed. Consider:

The formation of such hypothetical plans is something we engage in all the time—and with good reason. 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 am reading a biography of Napoleon, and on finishing a chapter discussing the details of the battle of Waterloo, 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 which involves just the sorts of commitments at issue here, albeit in the practical—as opposed to the epistemic—domain. (Schafer 2014b)

But what a weird thing to do. Why bother forming a plan for what I would do in Napoleon's shoes? And what role does this hypothetical plan play in my actual formations of plans for situations I will encounter? It's hard to see. We can see why we would want reasons to be universal if the role of such hypothetical planning was to inform our actual decision; but if not, it's hard to see why we should maintain this in any seriousness. Perhaps when we consider what we would do at the battle of Waterloo, we are entertaining what *we*, i.e. people with our values, would plan to do in such a situation. And this does not bear any direct relationship on the notion of an ontic reason that we have set out to justify.

On the other hand, it is reasonable to think that we might evaluate Napoleon's situation evaluatively. After all, it would be strange to find ourselves in a similar situation, but by no means impossible. And it is useful to have pre-planned responses to cases we might find ourselves in. Now, to be sure, it would make more sense for a general to engage in an evaluative judgment about Napoleon's reasons than it would for most of us, but the point is the same. It is reasonable to evaluate what *we*—people with our normative sensibilities—would do in this situation. But note that this does not give us much reason to think that ontic evaluation of Napoleon's situation is anything more than a curiosity. After all, what would the fact that Napoleon has ontic reason to sacrifice his men tell us about what we would have reason to do in a similar situation?

Continuing on this thread, we should wonder why we can so easily make sense of the question of what we would do in Napoleon's (tiny) shoes. After all, even if we assume our normative outlook is similar, that does not mean that we're not shifting our desires, loves, hates, fears, and so on. And the weight of various other factors can be difficult to suss out in advance. For example, it is be difficult to see, in advance, how we might react if we cared less about our own life, or more about the life of someone else, than we in fact do. But, nevertheless, we easily pass judgment on such counterfactual cases, as noted above by Schafer. Why is it so natural and easy for us to do so?

Presumably, the answer to this is that our judgments about what we ought to do in Napoleon's shoes are informed by the gross majority of our normative sensibilities *and* the projected assumption that Napoleon, different as he may seem to be from us, nevertheless shares much of our way of interacting with the world. Although normative outlooks can shift fairly significantly from agent to agent, nevertheless, it is very plausible that there is broad and pervasive overlap between any conceivable normative outlook and broad and pervasive overlap in how we interact with the world. Maybe it seems okay to cross against the light and shoplift from large chain stores to me and not to you, but a total lack of empathy towards other humans and random face-stabbings seem basically *out*. And asking ourselves the question of what we would do in Napoleon's shoes presupposes this fact. If not, then it is difficult to see how we could even get started.

Something like this, I believe, underlies Philippa Foot's nice observation:

But it is 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. . .

One imagines the reason that they were not struck by this is that it would be relatively *unhuman* to have no loyalty and devotion to one's city, family, and friends as against invading forces; it seems literally inconceivable that this could be a pervasive feature of folks and, though we accept that not everyone shares such a set of attitudes, nevertheless looking through the eyes of those who do not seems rather difficult. Where even to start?

If this is right, then we can see why we might mistake the *evaluative*-ish question of what *we* would do in Napoleon's shoes for the *ontic* question of what Napoleon has reason to do. Given the presupposition that we and he enjoy broad overlap in sensibility and normative outlook, the questions come to the same. So it is no surprise that we mistake the evaluative reasons our entertainment and evaluation of hypothetical situations yields for the ontic reasons that actually obtain for the agent in that case. As was remarked in the opening, this point should remind us of Wittgenstein's Woodsellers and how difficult is to understand their way of counting. The point is not that we cannot make out the way they measure; that much is clear. The difficulty is in 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 (or think we know) about measuring. It is entirely alien to us.

But if this is right, then the regulative role of considering bizarre abstract hypothetical situations is much less useful for informing our actual behavior than it otherwise appears. Imagining the normative outlook of someone who is Future Tuesday Indifferent (i.e. someone who could care less about pain incurred on future tuesdays) or loves, more than anything, counting blades of grass in massive oceans of lawn, would be nigh impossible. But then, as we are in no reasonable danger of being in this position, the role of such considerations in informing our actual plans is negligible. Which means that the demand that our normative reasons be entirely uniform between the cases is likewise undermined. Which means we have lost a significant reason to demand that our reasons not be subjective.

To briefly sum up this section: I have given an argument, based in the functional role of normative judgment, for thinking that facts about our reasons—at least general such facts—should be necessary. I then argued that, on inspection, we only need evaluative necessity to satisfy this role. I then argued that it's easy to see why we might mistake ontic necessity for evaluative necessity in this case: our evaluation of hypothetical cases presupposes some degree of uniformity of moral outlook and affective states. But, given this last point, we do not even need full evaluative necessity in order to satisfy the functional role of normative judgment. Evaluative human necessity (writing $c \sim_h d$ for d is a context similar to ours in the conventions, psychologies, and behaviors of the agents within):

$$\langle c, n \rangle \models \Box_{he} R(a, \gamma, \phi) \text{ if and only if } \forall c' [a, \gamma \in c' \wedge c' \sim_h c \Rightarrow \langle c', n \rangle \models R(a, \gamma, \phi)]$$

will suffice for normative judgments to play the role in organizing our actual behavior that Gibbard and Schafer suggest.

8. Counternormative Counterfactuals

This last discussion leads naturally into the final argument for normative necessity. Again, like the Kantian intuition mentioned above, this argument relies on our intuition of the falsehood of certain counterfactuals. Intuitively, the following counterfactual seems false

If I approved of random cruelty, I would have reason to be cruel

as do its analogues for other subjectivist views

If local moral conventions suggested random cruelty, I would have reason to be cruel

If I would be open to blame for not being cruel, I would have reason to be cruel

Intuitively, we do not have reasons to be randomly cruel, whether or not we disapprove of it.³² However, since we have distinguished two senses of necessity, we can and should distinguish two senses of these conditionals. Given a normative counterfactual conditional $\phi \Box \rightarrow \psi$, we can distinguish the evaluative normative counterfactual conditional

$\langle c, n \rangle \models \phi \Box \rightarrow_e \psi$ iff in the closest context c' in which $\langle c', n \rangle \models \phi$, $\langle c', n \rangle \models \psi$

from the ontic normative counterfactual conditional

$\langle c, n \rangle \models \phi \Box \rightarrow_o \psi$ iff in the closest context c' in which $\langle c', n_{c'} \rangle \models \phi$, $\langle c', n \rangle \models \psi$

The three examples above are clearly false on most subjectivist views if they are taken to be evaluative counterfactual conditionals. So long as we do not have, as a part of our normative standpoint, that the attitudes or conventions of people matter for whether or not 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.³³ Is this enough to downplay the intuition that such conditionals are intuitively false? Speaking for myself, I think it ought to be. It is entirely unclear why we would think that such conditionals are obviously false in the ontic sense. But I fear that this will be insufficient to convince most people. Luckily, more can be said.³⁴

³²Modulo lightweight reasons that arise from distinct contexts, like the fact that cruelty would be funny or the fact that I've promised to be cruel. I am committed elsewhere to the view that such reasons exist, but are nearly always outweighed. See REDACTED.

³³I will use 'approval' and 'disapproval' here in my discussion, but what I say holds for all the subjectivist metanormative views I sketched above.

³⁴Note that we can define a notion of an ontic conditional which expresses something similar to humanly ontic necessity

$\langle c, n \rangle \models \phi \Box \rightarrow_{ho} \psi$ iff in the closest context c' where $c' \sim_h c$ and $\langle c', n_{c'} \rangle \models \phi$, $\langle c', n \rangle \models \psi$

For derivative reasons, many of the counterfactuals will actually be false. Consider the closest possible world in which kicking dogs is smiled upon by our normative standpoint. What would such a world be like? It's tempting to say that it would be nearly identical to our world, except that in that world people approved of kicking dogs. But this is far too hasty. Our disapproval of kicking dogs is sourced in our empathy towards animals, the fact that dogs often play a close and intimate role in our lives, the fact that kicking them causes pain and so on. These facts do not sit nicely with approval of kicking dogs. It is far more plausible that the nearest possible worlds in which kicking dogs is smiled upon is one where one of the above source facts is false. Perhaps dogs either have a slightly different physiological make-up—one in which they actually enjoy or, anyways, don't mind being kicked—or perhaps dogs in that world are typically wild, scavenging food and terrorizing small children. In either case, though admittedly more in the former case, it is by no means obvious that kicking dogs has the obviously distasteful flavor that it has for us. Now, suppose that we do not feel any empathy towards animals. At all. I submit that our intuitive judgment as to whether or not we have reasons to kick dogs in such a case is not probative. As I suggested above, such a case differs so wildly from our actual case that we cannot make reasonable judgments about what reasons we would have if we were to be residents of the case. Our intuitions about such a case should thus not be assumed to be probative.

Likewise with cases where what we are supposing is even less normal for folks like us. For example, try to imagine what it would be like in order to approve, quite generally, of random cruelty or to accept norms permitting random cruelty (without assuming some massively mistaken set of beliefs about the world.) So much would have to change for us to think of general approval of such things that it is difficult to think about what the rest of such a scenario would be like—this is the Stroud/Wittgenstein point from above. Of course, the thought of being like that it morally repulsive, but note that this is an evaluative reaction, not an ontic one—perhaps it could likewise be interpreted as an ontic evaluation that presumes

That is, in less formal language, in the closest context c' where (a) c' is humanly similar to our context, and (b) ϕ holds, ψ holds. Something like “if we approved of kicking dogs, then we would have reason to kick dogs” will then turn out, plausibly, to be true since the closest world at which we approve of kicking dogs, where the agents are relatively like us, is one where kicking dogs doesn't cause them pain, but rather pleasure. Unfortunately, “if we approved of random cruelty, then we'd have reason to be cruel” also turns out to be true, though vacuously so, since there are no worlds where we approve of random cruelty, but where the agents are relatively like us. Such a world would be one where the agents are nothing like us.

overlap in normative standpoint without any evidence for such.³⁵ So, if we reach a clear reaction about the case, there are always two alternatives to thinking our reaction probative as to the truth or falsity of the relevant count. First, it might be an evaluative reaction we mistake for an ontic one. Second, it might be an ontic reaction for 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, at best, such a case massively dissimilar from our own as it consists of agents with massively incoherent normative standpoints. Not having reason to be randomly cruel, after all, is more or less a normative fixed point for us (Cuneo and Shafer-Landau 2014). As McCord and Cuneo remark, denying such a thing tends to provoke bewilderment (Cuneo and Shafer-Landau 2014, pp. 407-408) The case where they have coherent, but grotesque normative standards coherent with their approval of random cruelty is much more similar to our own since it doesn't involve denying that they agree with us on so much, but are so seemingly deviant in this one particular. And, in such a case, there is no obvious ontic reason for them to not be wantonly cruel, so the counterfactual is straightforwardly true.

At this point, it might be worried that there are real life cases of repulsive moral behavior: FGM, homophobia and sexism, actual and de facto slavery, overly severe retributive punishment for mild peccadillos, etc. It cannot be denied that these occur and it cannot be denied that such things are morally outrageous. Of course, the latter is an evaluative reaction. But, in sympathy with the worry, it is quite a pill to swallow that folks committing such moral crimes actually have reason to commit them and no reason (or, anyways, outweighed reason) to rescind. What can be said about this kind of worry? First and foremost, we should not neglect the degree to which many of these behaviors are sourced in mistaken or pernicious beliefs that such folks have reason to reject. For example, homophobia might be sourced in deep-seated religious beliefs of a nasty stripe, a mistaken understanding of, say, the role of religious tolerance, and the like. Given such beliefs, it is by no means obvious that someone's reasons point towards tolerance of those with different sexual preferences than those they accept; given the role of religious belief in their lives, they have reason to not act out of any empathy they have towards gay folks (as people). Likewise, there are reasons to think that many of our regrettable biases are sourced in mistaken empirical beliefs about people of a certain gender or race. Consider, for example, the pernicious effect of texts like

³⁵Note that this would mean that the evaluative and the ontic evaluation will more or less overlap, as mentioned above.

The Bell Curve.³⁶ The issue is complicated, of course, and I cannot do much here other than sketch the general outlines of this sort of approach: find, as much as possible, sources for immoral behavior in mistaken beliefs or mistaken derivative reasons of some type.

This, I think, is actually the sort of thing we *ought* to do when interpreting others. We should try as hard as possible to see them as not making moral errors, but factual ones. When we cannot find such errors, we should look for swamping beliefs—such as religious convictions that cut against human empathy—or mistakes of reason arising from mistaken application of beliefs to norms. 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 think it is, then there is significant reason to think that we generally presuppose that others are relatively like us. Otherwise there would be no reason to look so hard to explain away the immoral actions of others in terms of some mistaken or pernicious beliefs. And, when we cannot do so, it seems to me that we are in a position to them as we are to people who approve of random cruelty. Note how difficult it is to understand how someone could willingly participate in FGM or throw adulterers off of high towers. Their normative outlook is not just different from ours; it is 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 am not alone in this.

In this section, I have argued against the idea that the intuitive falsehood of counternormative conditionals gives us reason to believe in ontic necessity. I argued, in particular, that if we read these counternormative conditionals evaluatively, then they are uncontroversially false and that if we read them ontically, then even though they might be true, we should not rush to label this fact as counterintuitive. This for two reasons. First, because in many cases the closest possible world in which we, say, approve of kicking dogs, this can be explained by some difference in what our judgment about kicking dogs is sourced in. So, for example, it might be that the closest world in which we approve of kicking dogs is one

³⁶This is not to say that such issues aren’t very complicated. They are exactly that and more. See (Manne ms), for example, for a critique of the idea that we always treat people morally badly on grounds that they are less than fully human. As Manne notes, this plausibly plays some role, some of the time, but it does not account for all the ways in which we treat others in a morally grotesque fashion.

where kicking dogs is not painful to them. And, second, with more fundamental divergence in normative outlook—say, approving of random cruelty—the difference between their moral outlook and ours is so deep that we should not take our kneejerk reactions that the counternormative conditional is false as probative. It is more plausible that in this case we are mistakenly importing too much of our own normative standpoint into our evaluation of the reasons they would have from inside their normative standpoint.

9. Conclusion

I have argued against a number of considerations that are meant to support the idea that normative facts are necessary. In particular, I have distinguished and defined two separate notions of normative necessity—ontic and evaluative necessity—and argued that we can use either evaluative necessity or a restricted version of ontic necessity to do the work that full ontic necessity was meant to do. Evaluative necessity suffices to make sense of evaluating the reasons of others, significant overlap in our reasons suffices to make sense of normative communication when it can be made sense of (so, in particular, we can get away with presupposing overlap in normative standpoint), evaluating hypothetical cases requires an even weaker version of evaluative necessity—human evaluative necessity, and the resulting types of counterfactuals are unproblematic—evaluative counternormative counterfactuals are false, as should be expected, ontic counternormative counterfactuals are either false, unworrisome, or our intuitions about them are not probative.

I have relied in this defense on a number of claims about how stable and universal some of our normative standpoints are presumed to be. While I think these claims are very plausible, it may be that we should revise this presumption if it turns out that subjectivism is the correct view of normative matters. After all, even though 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. But this, note, is itself a normative claim and one which is not obviously true. Likewise, it is not obvious that we should revise our tendency to accept some form of evaluative normative necessity, just as it is not obvious that we should revise our tendencies to accept some form of evaluative aesthetic necessity. These tendencies play a useful prudential role in facilitating coordination, one which would be hampered by being

overly tolerant of the viewpoints of others.³⁷

On balance, it seems that the assumption that normative facts are necessary is on shaky footing. If there's an argument that favors it, it will not come from any of those we've considered. And it seems to me that without begging questions against a subjectivist point of view, it is unlikely to come from anywhere. This is not to say that there aren't other reasons to reject subjectivism; of course there may be good reasons to do so. But the putative necessity of normative facts—at least the putative *ontic* necessity 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 normative matters

References

- Ayer, A. J. 1946. *Language, Truth and Logic*. Victor Gollancz Publishers.
- Blackburn, S. 1988. Supervenience Revisited. *Essays on Moral Realism*: 59–74.
- Blackburn, S. 1998. *Ruling Passions*. Oxford: Oxford University Press.
- Blackburn, S. 2006. Must We Weep for Sentimentalism? *Contemporary debates in moral theory*: 144–159.
- Carballo, A. P. and P. Santorio. forthcoming. Communication for Expressivists. *Ethics*.
- Cuneo, T. and R. Shafer-Landau. 2014. The Moral Fixed Points: New Directions for Moral Nonnaturalism. *Philosophical Studies* 171(3): 399–443.
- Dreier, J. 1990. Internalism and Speaker Relativism. *Ethics*: 6–26.
- Dworkin, R. 1996. Objectivity and Truth: You'd better believe it. *Philosophy & Public Affairs* 25(2): 87–139.
- Einheuser, I. 2006. Counterconventional Conditionals. *Philosophical Studies* 127(3): 459–482.
- Foot, P. 1972. Morality as a System of Hypothetical Imperatives. *The Philosophical Review*: 305–316.

³⁷For suggestions in this direction, see (Woods forthcoming) and (Maguire and Woods manuscript). The sort of two-level explanation we offer there can be used to explain why we might have good reasons to engage in a practice, such as blaming others for not conforming to *our* aesthetic or moral standards, even though people's standards differ quite dramatically.

- Harman, G. 1975. Moral Relativism Defended. *The Philosophical Review*: 3–22.
- Kant, I. 2002. *Groundwork for the Metaphysics of Morals*. Yale University Press.
- Lenman, J. 1999. Michael Smith and the daleks: reason, morality, and contingency. *Utilitas* 11(2): 164–177.
- Manne, K. 2014. Internalism about Reasons: Sad but True? *Philosophical Studies* 167(1): 89–117.
- Manne, K. manuscript. Humanism: A critique.
- Martin, J. and G. S. Stent. 1990. I think; therefore I thank: A philosophy of etiquette. *The American Scholar*: 237–254.
- Parfit, D. 2001. Rationality and Reasons. *Exploring Practical Philosophy: From Action to Values*: 17–39.
- Peacocke, C. 2003. *The Realm of Reason*. Oxford University Press.
- Rosen, G. What is Normative Necessity?
- Schafer, K. 2014a. Constructivism and Three Forms of Perspective-Dependence in Metaethics. *Philosophy and Phenomenological Research* 89(1): 68–101.
- Schafer, K. 2014b. Doxastic Planning and Epistemic Internalism. *Synthese* 191(12): 2571–2591.
- Schroeder, M. 2007. *Slaves of the Passions*. Oxford University Press.
- Smith, M. 1994. *The Moral Problem*. Blackwell.
- Smith, M. 1999. The Non-arbitrariness of Reasons: Reply to Lenman. *Utilitas* 11(02): 178–193.
- Sobel, D. 2009. Subjectivism and Idealization. *Ethics* 119(2): 336–352.
- Sobel, D. forthcoming. Subjectivism and Reasons to be Moral. *From Valuing to Value*.
- Stalnaker, R. 2013. Assertion. *The Semantics-Pragmatics Boundary in Philosophy*: 179.
- Street, S. 2009. In Defense of Future Tuesday Indifference: Ideally Coherent Eccentrics and the Contingency of What Matters. *Philosophical Issues* 19(1): 273–298.

- Street, S. 2010. What is Constructivism in Ethics and Metaethics? *Philosophy Compass* 5(5): 363–384.
- Street, S. 2012. Coming to terms with contingency: Humean constructivism about practical reason. *Constructivism in Practical Philosophy*: 40–59.
- Väyrynen, P. The Supervenience Objection to Moral Non-Naturalism.
- Velleman, J. D. 2013. *Foundations for Moral Relativism*. Open Book Publishers.
- Woods, J. forthcoming. The Normative Force of Promises. *Oxford Studies in Normative Ethics*.