Epistemic Consent and Doxastic Justification[1]

Luis Oliveira
University of Houston

My starting point is what we can call the *Normative Authority Conception* of justification, where S is justified in their belief that p at t (to some degree n) if and only if their believing that p at t is not ruled out by epistemic norms that have normative authority over S at t. With this in mind, this paper develops an account of doxastic justification by first developing an account of the normative authority of epistemic norms. Drawing from work in political philosophy, I argue that (a) the cognitive and evaluative commitments and concerns behind our actual practices of holding each other and ourselves accountable for our beliefs reveal which epistemic norms we have *consented to be under*, and that (b) it is because we have consented to be under the authority of these norms – by actually holding ourselves and others accountable to them – that they in turn have normative authority over us. By connecting the authority of epistemic norms to the authority we have over ourselves in this way, the resulting account of doxastic justification (i) explains why it can be appropriate to criticize, resent, or sanction someone for having unjustified beliefs, (ii) avoids the phenomena of normative alienation and normative parochialism, and (iii) respects the social and collective nature of epistemic justification.

Introduction

In various recognizable ways, we hold ourselves and each other accountable for the things that we believe. Though not typically obsessive about it, most of us sometimes reflect on our own beliefs and cognitive powers, especially when under pressure from other beliefs, sensations, feelings, testimony, disagreement, and so on. Most of us can tell when something seems "epistemically off" – though not always and not always correctly – and most of us are at least somewhat disposed, in varying degrees of seriousness, to do something about that bothersome impression.[2] No doubt we are industrious in communicating any perceived epistemic foibles in our peers. What others believe and how they manage their cognitive lives, especially on sensitive topics or topics near and dear to us, is often the very target of our criticism, the cause of our resentment, and the grounds for sanctions such as social exclusion and public shaming.

My starting point in this paper is a commitment to understanding the property of "being doxastically justified in believing that p" as the inchoate ideal that is behind these familiar (but hard to describe) *practices of epistemic accountability*. When we try to understand the proper grounds for evaluating

[2] See Zagzebski (2012, 27-33) for a discussion of the epistemic fundamentality of a pre-reflective "experience of dissonance".

someone's state of believing that p as being justified or unjustified, I take it that we are trying to understand the proper grounds for engaging with ourselves and each other in ways that we are already disposed to engage. In my view, this gives a clear purpose and a clear value to this kind of philosophical investigation. We are beings who are inextricably engaged in, and dependent on, these private and collective practices of epistemic accountability; by better understanding and by improving on these practices, we better understand and improve ourselves.[3]

There are three dominant conceptions of doxastic justification in the literature: one focusing on whether the mechanisms responsible for the relevant belief are connected to truth in the right way, one focusing on whether that belief is based on adequate grounds, and one focusing on the relation between that belief and character traits of the believer. Consider three traditional and representative articulations:

> If S's belief in p at t results from a reliable cognitive process, and there is no reliable or conditionally reliable process available to S which, had it been used by S in addition to the process actually used, would have resulted in S's not believing p at t, then S's belief in p at t is [doxastically] justified. (Goldman 1979, 123)[4]

> S's doxastic attitude D at t toward proposition p is [doxastically justified] if and only if (i) having D toward p [fits the evidence S has at t]; and (ii) S has D toward p on the basis of some body of evidence e, such that (a) S has e as evidence at t; (b) having D toward p fits e; and (c) there is no more inclusive body of evidence e' had by S at t such that having D toward p does not fit e'. (Feldman and Conee 1985, 93)[5]

> A justified belief is what a person who is motivated by intellectual virtue, and who has the understanding of his cognitive situation a virtuous person would have, might believe in like circumstances. (Zagzebski 1996, 241)

Different accounts have been developed along similar or hybrid lines. Yet if these accounts – or any other – are about the property that *properly grounds* our practices of epistemic accountability (their original intentions notwithstanding), then they must be taken as more than simply the articulation of a concept, a point of view, a natural kind, or some such: they must be taken as expressing *epistemic*

---

[3] My commitment to relating doxastic justification to our practices of epistemic accountability does not reflect a sociological belief that epistemologists working on this topic are, by and large, also seeing things this way. See Plantinga (1990, 45-49), Alston (2005, 12-15), Silva (2017, 28-29), and Goldberg (2018, 13-47) for discussion of the variety of different conceptions of justification found in the literature, and for very different suggestions for what to do about it. Following some work on externalism about knowledge, moreover, I here take it that knowledge does not require justification (c.f., e.g., Armstrong 1973, Nozick 1981, Audi 1993, Sosa 2007, Pritchard 2005, Audi 2020). I also do not assume that *justification* and *rationality* amount to the same thing (cf. Siscoe 2021).

[4] I have inserted 'doxastically' to clarify the kind of justification at stake. Goldman's preferred term for this property, a term he introduces soon after the quote just above, is "*ex post* justification" (cf. Goldman 1979, 124).

[5] I am substituting 'doxastically justified' for Feldman and Conee's preferred term for the same property, 'well founded', and I am inserting 'fits the evidence S has at t' in place of their place holder use of 'justified', which they define earlier as fitting the evidence (cf. Feldman and Conee 1985, 83).

norms as well, norms about the formation and maintenance of doxastic states, norms that we can hold and be held accountable by. In particular, we must be able to derive from them true claims of the form "S ought not ϕ at t," claims whose truth explains why it can be appropriate to criticize, resent, and sanction those who do, regretfully, ϕ at t. Even more, given the various kinds of norms we are ordinarily familiar with – moral norms, prudential norms, conversational norms, norms of etiquette, and so on – the epistemic norms derived from these and other accounts of doxastic justification must have *normative authority* over us: it must be truly appropriate to criticize, resent, or sanction someone for flouting *them*.[6]

What I have as my guide in this paper, therefore, is what we can call the *Normative Authority Conception* of justification:

> **NAC:** S is justified in their belief that p at t (to some degree n) if and only if their believing that p at t is not ruled out by epistemic norms that have normative authority over S at t.

This is not a particular *account* of doxastic justification, one that by itself rivals the traditional views just mentioned. This is instead a guiding *conception* indicating what accounts of doxastic justification are supposed to be about and, therefore, a guide for how to judge their plausibility.[7] We get particular accounts of justification out of this conception, of course, by endorsing an account of normative authority – an account of what makes the violation of a norm in general, or an epistemic norm in particular, the proper ground for criticism, resentment, and sanctions. No substantive account of justification is therefore eliminated by *fiat* on this conception, even if accepting it imposes substantive and controversial constraints on what could count as a plausible candidate.

Pleasant ecumenicism aside, it is very much unclear to me whether the epistemic norms expressed by the accounts just above, and the particular prescriptions we can derive from them, have normative authority over us. It is unclear to me, that is, whether epistemic norms enjoining us not to have unreliably formed beliefs, improperly based beliefs, or unvirtuous beliefs, can properly ground our practices of epistemic accountability. That said, I will not argue for this claim today. Instead, with this concern in mind, in this paper I develop an alternative account of doxastic justification by looking

---

[6] See Brown (2019, 2020) for a similar claim about "epistemic blameworthiness." We need not assume, however, that epistemic norms have *independent* normative authority, so long as their authority can be derived from norms that do, indeed, have normative authority over us. I won't discuss this kind of indirect approach here. That said, see Oliveira (2017) for a criticism of accounts connecting epistemic norms to moral norms (cf., e.g., Grimm 2009), and see Cote-Bouchard (2015) for a criticism of accounts connecting epistemic norms to prudential norms (cf., e.g., Kornblith 1993).

[7] See Alston (1985) for the distinction between conceptions and accounts and for the theoretical priority of the former. See Silva (2017) for discussion and application.

for an explanation for the normative authority of epistemic norms – for why flouting them can make criticism, resentment, and sanctions appropriate. In a reversal that has Strawsonian and Alstonian connections, I will argue that our practices of epistemic accountability have explanatory priority over the grounds of their own propriety.[8]

Drawing from work in political philosophy, I will argue that (a) the cognitive and evaluative commitments and concerns behind our actual practices of holding each other and ourselves accountable for our beliefs reveal which epistemic norms we have *consented to be under*, and that (b) it is because we have consented to be under the authority of these norms – by actually holding ourselves and others accountable to them – that they in turn indeed have normative authority over us. In other words, it is by doing well or poorly with respect to our own epistemic commitments, collectively and interpersonally thus understood, that we are justified or unjustified in our beliefs and, consequently, properly subject to criticism, resentment, and sanctions for them. What I offer here, to be clear, is but a sketch. My aim is to identify a particular difficulty, motivate a potential solution, and outline a program for its further development.

1. Attributability and Ownership

I want to begin by quickly surveying some work that is motivated by a similar concern to mine. The concern here is with what makes our practices of epistemic accountability appropriate. I will argue, however, that there are two dimensions to this question of propriety and that these views only address one of them – successfully or not.

Working from within the Feldman and Conee general approach, Adam Leite has argued that we can only be the proper targets of criticism for our beliefs if we have the power to directly determine, through private or public deliberation, the reasons for which we hold our beliefs. This is because our reasons for believing that p – which are the target of criticism, on views like Feldman and Conee's – must be *attributable* to us, and because "states of a sort which are never directly determined by a person's conscious deliberation, the commitments incurred through her conscious deliberation, or her best explicit evaluation of reasons, are not attributable to her in the relevant sense" (Leite 2004, 231). This concern with attributability as a requirement for proper criticism is certainly not unique to Leite. The kinds of evaluations behind our ordinary practices of accountability – epistemic or not –

---

[8] See Strawson (1962) and Alston (1989, 1991). For some difficulties for this explanatory reversal regarding "moral responsibility," see Todd (2016). For a way out of those difficulties – a way out that is available for my proposal here about "doxastic justification" – see Beglin (2018).

4

seem to presuppose that we are, in some relevant sense, the sources of whatever it is we are being criticized, resented, or sanctioned for. It would certainly be odd to behave in this way towards someone in virtue of something that merely happened to them.[9] Leite's view on what it takes for a state to be attributable to someone, however, where one must be involved in *directly determining* that state, is indeed somewhat unique.

Yet if we take the Feldman and Conee approach to doxastic justification, and if we think of justification as the inchoate ideal behind our practices of epistemic accountability, and if we moreover accept such a tight connection between direct determination and attributability, we are then led to something like the view that q is S's reason for believing that p only if S "endorses its adequacy as a reason for believing that p, and is committed to responding in appropriate ways if q proves to be an inadequate reason" (Leite 2004, 237). Failing this condition, then, whatever else is true of q, it is not my reason for believing that p. This is not, of course, an account of what it takes for something to be a *good* reason, nor is it an account of what it takes for S to be justified in believing that p. It is simply an account of one's *actual* reasons, one that follows from substantive views about what it takes for someone to be a proper target of criticism by virtue of their reasons for believing that p.[10]

In defending some such view, as you might expect, Leite is not suggesting that we only have reasons for our beliefs and are therefore capable of being justified in holding them or criticized in the relevant ways, after having explicitly committed ourselves to them in private or public deliberation. Even though criticizing someone as unjustified is, for Leite (2004, 233), to criticize them for their "*endorsement* of certain considerations as adequate reasons" (my emphasis), that endorsement can be tacit and dispositional. One counts as endorsing q as their reason for p, according to Leite (2004, 234), if one *would* "appropriately manifest, express, or acknowledge one's reasons in one's explicit thinking or attempts to justify one's belief," *were* the appropriate situations to present themselves. This is all well and good; we would have implausibly too few justified beliefs otherwise.

Less stringent accounts of what it takes for a state to be attributable to someone are friendlier to Goldman's approach to justification. On this approach, recall, justification is not determined by one's reasons for believing some p but rather by properties of the mechanisms that produced the belief, mechanisms that are sub-personal, typically involuntary, and almost always opaque to introspection.

---

[9] For broader discussions of the connection between normative evaluation and attributability see Scanlon (1998), Watson (2004), and Smith (2012).
[10] For an evaluation and critique of Leite's full epistemological project, see Perrine (2018).

Nonetheless, the states that are outputs of these mechanisms can be attributable to us – in the sense that makes us proper targets of normative evaluation for having them – if we *take ownership* for those mechanisms and those outputs. If we identify these sub-personal mechanisms with ourselves, and identify their outputs as our own, then we are properly criticizable when these outputs don't look good.[11]

There are at least two kinds of proposals regarding this ownership-relation in the literature.[12] One proposal takes ownership for a belief as primarily a matter of recognizing oneself, through a historical and developmental process, as the source of beliefs formed by that mechanism, and as a matter of accepting that one is a fair target of criticism for having those beliefs (cf. Breyer and Greco 2008; Breyer 2010, 143; McCormick 2011, 173-4). The other proposal takes ownership for a belief as a matter of one's reflective attitude over one's reasons: either the attitude of taking oneself to have better reasons for endorsing that belief than not (cf. Duke-Yonge 2013, 246), or the attitude of recognizing the judgment that one's reasons support that belief as reflecting one's values (cf. Osborne 2021). Once again, these are not accounts of what it takes to be justified in believing that p, but rather accounts of what it takes to own up to one's beliefs – of what it takes for the belief that p to be attributable to S, in the relevant sense – such that one can be a proper target of criticism on account of having them.

None of these accounts, moreover, requires an explicit and conscious act behind the fact of ownership. From the mechanism side, McCormick (2015, 121) tells us that "even if we never consciously endorse a mechanism, we can still have ownership of it… if *my practices* reveal that I have accepted the expectation that I keep beliefs of this kind in line with my higher order judgments of how I ought to believe" (my emphasis). From the reasons side, Osborne (2021, 8) tells that S owns a belief "just in case she holds it for reasons she takes *or is disposed* to take herself to possess, and the attitude reflects an evaluative judgment that she regards *or is disposed to regard* as her own, i.e., *she is disposed* to reflectively endorse it as expressing her values" (my emphasis). Once again, it is hard to see how this could be otherwise, our practices of epistemic accountability being so profligate, and our time spent on actual discussion and actual deliberation being so limited.

---

[11] For a broader discussion of the connection between normative evaluation and ownership, see Fischer and Ravizza (1998).
[12] Some have discussed belief ownership in connection with doxastic justification (Breyer and Greco 2008, Breyer 2010, Duke-Yonge 2013), others in connection with doxastic responsibility (McCormick 2011, Osborne 2021). For the purposes of this paper, I am ignoring these differences. See Oliveira (2018), however, for an argument that epistemic justification should not be understood as an evaluation of an agent as responsible in their belief.

Despite their various differences, what all of these views have in common is their focus on the conditions required for S to be a *proper target* of criticism for believing that p, in particular the requirement that S's belief be attributable to S. In this way, however, all of these views focus on what it is *on the subject's side* that makes doxastic evaluations appropriate. What they leave out, in turn, is a discussion of the conditions required for *a norm* to be an appropriate standard for the evaluation of that subject's state. What gives normative authority, after all, to certain norms as the proper standards for evaluating the beliefs and/or reasons that are properly attributable to me? In owning up to certain beliefs, or in committing to certain reasons, what makes me accountable for flouting *these* epistemic norms and not others? There are two dimensions, in other words, that together constitute the propriety of our practices of epistemic accountability. What is being evaluated positively or negatively must be properly attributable to me, on the one hand, and the grounds for that evaluation must have authority over me, on the other. I can define a host of properties and evaluate you on whether you instantiate them, but that won't automatically make it appropriate for me to criticize, resent, or sanction you for falling short, even if they are attributable to, or owned by, you.[13]

Consider an analogy. When Magnus Carlsen evaluates a complicated chessboard position as "winning for black" – where this does not mean, of course, that black will in fact win – the evaluation is appropriate because Carlsen is both (a) evaluating a game of chess and (b) applying to it the appropriate standards for that kind of evaluation. It would be quite irrelevant if Carlsen were to evaluate a game of backgammon by those same standards, or were to evaluate a game of chess by the standards appropriate to backgammon instead. Appropriate evaluations require a match between proper targets and proper standards. And when those standards are behind practices of accountability, then an account of what makes them *proper* just is an account of what makes gives them *normative authority*. Without some such account, however, our understanding of our practices of epistemic accountability, and our search for an account of doxastic justification that can ground them, is incomplete.[14]

2. Constitutivism and Normativism

---

[13] To be clear, none of these authors is trying to answer the question I'm pursuing here, so I am not imputing their arguments with this failure. My claim is that even if we take them as having achieved their aims – i.e., take them as having correctly identified the conditions under which we are the proper targets of evaluation for having the belief that p – we are still in the dark as to which norms are such that violating them makes criticism, resentment, and sanctions appropriate.
[14] Discussions of normativity and attributability in moral theory typically ignore the issue I'm highlighting for a very natural reason: there is a strong presumption, in the moral case, that the relevant norms are indeed authoritative in the sense of properly grounding criticism, resentment, and sanctions. I have two comments here. First, whatever we say about the appropriateness of this presumption in the moral case, there is no similarly strong presumption in the epistemic case. Second, I'm of the view that such presumptions, of whatever strengths, must ultimately be backed up by arguments anyway.

I now want to consider a kind of account of epistemic normativity that, by focusing on the source of epistemic norms, "promises to explain, in a relatively straightforward way, the authority of epistemic norms" (Nolfi 2015, 181). On this kind of account, "beliefs will be appropriately evaluated with respect to epistemic norms simply by virtue of being the sorts of mental attitudes that they are" (ibid). Despite their recent wide support, however, I will argue that these accounts fail to deliver on that promise. I begin by distinguishing between two minor variants of this view.

One variant is often called *constitutivism*. The general suggestion here is that we can extract norms for engaging in any kind of activity by examining the non-normative constraints on what counts as performing that activity in the first place. One only counts as playing chess, for example, if one is moving certain kinds of pieces in certain kinds of ways. This is to say that it is constitutive of the activity of playing chess – as opposed to playing some other game or no game at all – that one refrains from doing certain things at certain times (e.g., from moving one's rook diagonally during a match). These constitutive constraints are non-normative in the sense that they are merely claims about certain actions and certain activities, and not yet claims about norms. Nonetheless, from these claims about what counts as chess playing, we can *indirectly* extract norms of the kind "if S is playing chess at t, then S ought not Φ at t," where Φ-ing at t is incompatible with counting as playing chess at t.

We can derive epistemic norms in a similar way. The classic, if imprecise, statement of this view is Bernard Williams' (1973, 136) famous claim that "beliefs aim at truth," which is intended as a non-normative description of the role a mental state must actually play in one's psychology in order to count as a belief in the first place. Statements along these lines abound. According to Scott-Kakures (1993, 87), "no one believes that p if she also believes that the belief that p is unsupported by any consideration having to do with the truth of p." According to Street (2009, 225), "S believes that p if and only if (as a routine matter or when she thinks about the matter in full consciousness) S takes considerations that S regards as bearing on the truth of p to be reasons for and against believing that p." These are not simply (alleged) empirical observations or predictions; these are substantive non-normative constraints on what kinds of mental states can count as a belief. They are analogous to the claim that "no one is playing chess if they are moving their rooks diagonally on the board." Yet, from such constraints on what counts as a belief, we can indirectly identify epistemic norms of the kind "if S is believing at t, then S ought not Φ at t," where Φ-ing is incompatible with counting as believing

that p at t. Identifying a plausible and satisfying non-normative constraint on the mental state of belief, of course, has proved to be tricky business.[15]

Another variant of this kind of account is often called *normativism*. Here the suggestion is not that we derive epistemic norms indirectly by first identifying the non-normative constraints on what count as a belief, but rather that we do so directly: it is part of the very nature of the mental state of belief that it is subject to certain norms. On one development of this view, for example, our concept of belief has a constitutive standard of correctness, namely, that "S's belief that p is correct iff p is true" (cf. Shah and Velleman 2005). On another development, the mental state of believing that p "is correct if and only if it is disposed to inform our actions by serving as a kind of map so that our actions successfully achieve the ends that our actions are meant to achieve" (Nolfi 2015, 197-8). Despite the differences between these two developments – one focusing on our concept and one focusing on a mental state – and despite the differences between this variant and the previous – one being indirect and one being direct – the central move is nonetheless the same. From the very nature of belief, we can extract norms of the kind "if S is believing at t, then S ought not Φ at t," this time where not being subject to this norm is itself incompatible with counting as believing that p at t. Of course, arriving at a plausible and satisfying formulation of this norm has proved to be tricky business just as well.[16]

For my present purposes, it matters very little which variant, or which development of which variant, is correct. Constitutive normativity of whatever kind is rather cheap. The norms extracted from the nature of chess playing (directly or indirectly) don't really carry, by themselves, any normative authority over us. I can at any time decide to stop playing chess and make whatever moves I wish, or no moves at all, and – provided I have not thereby flouted any *other* norms of, say, respect, or promise keeping, and so on – no kind of criticism of my behavior would be appropriate on those grounds. The constitutive norms of chess may well properly ground evaluations of chess-playing, but those evaluations don't by themselves properly ground criticism, resentment, and sanctions. To be sure, epistemic norms are different from the norms of chess in that we have no way of begging off from the activity of believing, no way to decide not to form or hold beliefs, and so no way to remove ourselves from the domain of application of the norms governing that kind of activity. Perhaps, then, this difference grants them the normative authority lacked by the norms of chess and the norms of similarly contingent activities.

---

[15] For relevant discussions, see Winters (1979), Bennett (1990), Scott-Kakures (1993), and Frankish (2007).

[16] For relevant discussions, see Shah and Velleman (2005), Hieronymi (2008), McHugh and Whiting (2014), and Nolfi (2015).

But this is to misunderstand the problem. From the fact that I cannot escape a certain activity, it doesn't follow that the norms governing that activity have normative authority over me – that it is proper to criticize, resent, and sanction me for flouting them. This is because the following generalization is not true: for all activities ψ, if S is engaged in ψ-ing then S ought to ψ well. If this were true and I truly couldn't help but ψ, then it would be inescapably true of me that I ought to ψ well. And since flouting the constitutive norms of ψ-ing counts as a particularly egregious way of not ψ-ing well – either by not ψ-ing at all or by necessarily ψ-ing incorrectly – it would follow from all of this that those norms have normative authority over me, that I am inescapably *bound* by them in the sense capable of grounding proper criticism, resentment, and sanctions. I ought to ψ well if I ψ at all; I ψ-ed badly; shame on me. But the relevant generalization is not true. From the fact that I am engaging in a certain activity, it may follow that my performance is appropriately evaluable by the norms of that activity, but it does not follow that those norms have normative authority over me and that I can be properly held accountable for my shortcomings. So if the norms of belief have normative authority over us, it is not *because* they are the constitutive norms of something, even if it is something we can't escape.[17]

This worry applies to any bare appeal to constitution, no matter how involved. Recently, for example, A.K. Flowerree (2018) has defended a view that locates the source of epistemic normativity not in the nature of belief but rather in the nature of agency. Following Korsgaard's (2008) account of action, Flowerree (2018, 305) suggests that "our understanding of an action is intelligible only if we see it as the agent engaging in movement that is *guided* by her representation of the world. So any description of what the agent is doing will have to include what the agent took to be true when she acted. It is part of what makes the action the action it is" (emphasis original).[18] With belief and action thus intertwined, the constitutive norms of belief can inherit the normative authority of the

---

[17] Osborne (2021, 9) says that "the holding of a belief by default makes one answerable for that belief. The intuitive justification for this claim is that if someone e.g., professes to believe *that* p, it always is apt (if not *appropriate*) to ask them *why*" (emphasis original). But the question, once again, is not whether the norm applies; the question is whether it has normative authority. It is always appropriate to evaluate a chess move by the rules of chess, but this doesn't make those rules authoritative over me (in the relevant sense of accountability). See Enoch (2006, 2011) for a related and influential criticism of constitutive accounts of "agency". The key difference between the challenge Enoch presents and mine is that mine is not looking for what reasons one has to *care* about constitutive norms. See McHugh and Whiting (2014, 707-8) for a discussion of this challenge in terms of hypothetical vs instrumental norms. My challenge is orthogonal to that distinction. See Perrine (forthcoming) for discussion of similar issues regarding *epistemic value*.

[18] There is an echo here of Nolfi's (2015, 197) claim that "the proper function of belief is to inform our decisions to act by serving as a kind of map of the way things are so that we achieve whichever ends our actions aim to achieve," though the views they develop differ in important ways.

constitutive norms of agency. As Flowerree (2018, 305-6) puts it, "believing well is constitutive of acting well."

The success of this account of epistemic normativity, of course, depends on the success of the supporting constitutive account of practical normativity. The constitutive norms of belief have authority over us, on this picture, only if the constitutive norms of agency have authority over us *by virtue of being constitutive of agency* – and not on the basis of something else. Yet constitutive accounts of practical normativity – those that extract the norms of action from what is constitutive of agency – face precisely the challenge outlined above. What gives the constitutive norms of agency *their* normative authority? The answer cannot be an appeal to the mere fact that they are *the constitutive norms* of agency. We have already seen that it is false that, for all activities ψ, if S is engaged in ψ-ing then S ought to ψ well. What is constitutive of some activity, by itself, is neither here nor there for the matter of normative authority. Even if Flowerree is correct that "believing well is constitutive of acting well," we have not yet been told why it is the case that we ought to "act well," such that it would follow that we ought to "believe well" too. Whatever gives the constitutive norms of agency – if indeed there are any some such – their normative authority, such that they can pass it along to the norms of belief, must be some further fact not yet described.

Neither constitutivism nor normativism, therefore, delivers on the promise "to explain, in a relatively straightforward way, the authority of epistemic norms" (Nolfi 2015, 181). To do so, they must either insist, implausibly, on the generalized claim that, for all activities ψ, if S is engaged in ψ-ing then S ought to ψ well, or identify some fact other than the fact of constitution that explains why we ought to ψ well in this or that case. Either way, we better look elsewhere for what we're after.[19]

3. Consent and Authority

I want to suggest that we can make progress towards an account of the normativity of justification by paying attention to work on political legitimacy instead.[20] There is a long tradition in political

---

[19] Constitutivists sometimes insist that questioning the authority of the constitutive norms of agency is incoherent (see Velleman 2009 and Flowerree 2018). Since *all* normative authority derives from the constitutive norms of agency, on their view, to ask for the source of *their* normative authority is to pose an ill-formed question. This strikes me as objectionably self-confident. Of course, if constitutivism is true, then the challenge is misguided; that's a simple logical entailment of the claim that, if the challenge is not misguided, constitutivism is false. But to say that there is no challenge in the first place, that its very statement is incoherent, precisely *because* of what is postulated by the theory being challenged, that's theoretical navel-gazing at its worst.

[20] Drawing general lessons about normativity – or particular lessons about epistemic normativity – from work in political philosophy is not unusual. See, e.g., Raz (1988) and Zagzbeski (2012).

philosophy behind the idea that a government's legitimate rule depends somehow on the consent of those being ruled by it. This connection between consent and legitimacy, as I see it, is largely parasitic on the fundamental intuition that only the act of consent can transfer the authority we have over ourselves to something or someone who does not have that authority of necessity. "The ultimate authority over oneself," I here agree with Zagzebski (2012, 136), "is oneself." What I want to suggest, therefore, is that the *normative authority* of epistemic norms is very much like the *legitimacy* of a government's rule over a people. Both properties arise when groups of people voluntarily organize their lives together under common rules, and both properties can attach themselves, at different times, to a varied and sometimes even incompatible set of rules. By examining the relation between consent and political legitimacy, at any rate, I hope to draw lessons about the connection between consent and normative authority as well. In this section, I describe Amanda Greene's work on political authority and I extract from it a general principle outlining a connection between consent and normative authority.

Connecting consent to a plausible theory of political legitimacy is no mean feat. While the *actual* consent of all individuals does not seem necessary for a government's legitimate rule over them, the *hypothetical* consent of all individuals does not seem sufficient for legitimate rule either. One of the obvious problems of requiring actual consent for legitimacy is that it would turn out that no government is, or has ever been, or can ever be expected to be, legitimate. This seems implausible.[21] But merely requiring hypothetical consent for legitimacy seems inadequate too. If an entire group of people objects to a particular government policy, it is not enough for a government official to reply – even if truly – that the policy is legitimate since they *would* consent to it, under ideal conditions of information, rationality, and evaluative orientation. Actual dissent can have consequences for legitimacy, even if against a policy that "all citizens as free and equal may reasonably be expected to endorse in the light of principles and ideals acceptable to their common human reason" (Rawls 1996, 137). Any plausible theory of how consent contributes to legitimacy, therefore, must somehow find a middle way.

Bearing in mind these worries, Amanda Greene (2016, 2017, 2019, 2020) has recently articulated and defended a "Weberian" account of political legitimacy, one where "a state is legitimate to the degree that it is regarded by its subjects as having a valid claim to rule" (2017, 314). The heart and burden of this kind of account, of course, is specifying the conditions under which such acknowledgment of

---

[21] Some, of course, accept this result as true (e.g., Huemer 2012).

validity – such consent[22] – truly contributes to legitimacy. In short, Greene takes any kind of authority to have an essential function and takes legitimacy as a status that some authority can have by virtue of being recognized as performing its essential function well. Consent that is disconnected from such *telic recognition*, as I will call it, fails to contribute to the attainment of this status.[23]

In the case of governments, the key characteristic of what Greene calls *quality* consent – consent that truly contributes to legitimacy – is that the consent is based on a positive evaluation of the government's success with respect to the essential function of governance: benefitting its subjects through the exercise of power and authority (cf. Greene 2016, 84; 2019, 69). This is the sense in which the kind of evaluative attitude behind quality consent has a *fixed target*. Consenting to the authority of a government on the basis of something that is unrelated to the essential function of governance does not contribute to that government's legitimacy since it fails the Weberian requirement of telic recognition – it fails to be a "recognition of a political order as a political order" (Greene 2019, 72). But the relevant evaluative attitude, while having a fixed target, also has a *flexible content*. As Greene puts this at different times:

> On the basis of what a subject considers relevant and valuable *by her own lights*, she forms an overall subjective assessment of governance and, on that basis, consents to her political order. (2016, 84, my emphasis)

> The judgments that are involved rely on *subjective valuations of apparent goods* that have been achieved by the regime, at either the individual or collective level. (Greene 2019, 70, my emphasis)

What is flexible here, to be clear, are the terms of the evaluation behind quality consent: what *I* consider beneficial in a government's exercise of power and authority. It is this combination of a fixed evaluative target with a flexible evaluative content, in an act of telic recognition, that gives quality consent its legitimacy-producing power.

Greene's neo-Weberian connection between political legitimacy and quality consent is attractive for a variety of reasons. I want to highlight three reasons in particular that will be presently relevant. First, Greene's account respects the fundamental intuition, mentioned above, that the authority of a government is somehow derived from the authority we have over ourselves, an authority that is

---

[22] In her 2016, Greene speaks in terms of "consent". In subsequent work, Greene speaks in terms of "assent". I will use "consent" throughout, however, except for direct quotes from her later work, since "assent" already features prominently in other debates in epistemology (e.g., Owen 2003, Williamson 2007, Baxter 2018).

[23] The term "recognition" here is factive. As Greene (2020, 214 fn 2) puts it, "legitimacy consists of an objective element (*fulfilling* a promise to be valuable) and a subjective element (being seen as such by the relevant audience)" (my emphasis).

guided and expressed by our actual, subjective, and variegated evaluative and cognitive states. In this way, her account correctly prevents the co-existence of legitimate authority and widespread alienation of the kind that arises when "one cannot authentically affirm the regime to which one is subject as in any sense worthwhile, i.e., as providing at least some benefits that one values subjectively" (Greene 2019, 78). Second, despite locating the ultimate source of political legitimacy in the individuals themselves, Greene's account nonetheless respects the social and collective nature of this status. A government does not fail to have legitimate authority over me simply because I fail to consent to it, even though my consent or dissent matters for its legitimacy. Third, Greene's connection between political legitimacy and quality consent avoids narrow-minded parochialism: it does not consider a state as illegitimate simply because it does not conform to the value judgments of my culture (e.g., with the fundamentality of personal freedom) or the political ideals of my time (e.g., with democratic liberalism); rather, it allows legitimacy to obtain despite variations in value systems (cf. Greene 2017, 315). If an illiberal and undemocratic state is benefitting its subjects through the exercise of power, and if its subjects recognize it as such and consent to its authority despite lacking a degree of personal freedom that we ourselves would find minimally adequate, then that state counts as legitimate by Greene's neo-Weberian lights. All of this seems right to me.[24]

Just as we have seen above with one's commitment to reasons or one's ownership of mechanisms and beliefs, however, actually consenting to a government's authority does not require either overt expressions with a specific affirmative content or similarly structured occurrent thoughts. Instead, one typically expresses one's consent in behavior that expresses the relevant evaluative and cognitive attitudes. "Normally," as Greene (2019, 71) puts it, "the disposition to comply and readily cooperate with exercises of power and authority is a good indicator of this acceptance [i.e., quality consent]." A bit later,

> Quality assent stands for a combination of a volitional state (acceptance) and a belief state (judgment). A decent proxy for this combination of willing and believing is a set of behavioral dispositions – to obey, first and foremost, but also to evince general support for the order upheld by the regime. (Greene 2019, 76)

The qualifiers "normally" and "decent" are unavoidable here since non-consenting subjects may have prudential reasons to comply and cooperate with their governments that are unrelated to consent. But what is important for our present purposes is the fact that our dispositions can instantiate our consent, and not so much the difficulties that beset illuminating the correct connection between the

---

[24] What I have presented here, of course, is not a detailed and careful defense of Greene's views. For a fuller picture, see Greene (2016, 2017, 2019, 2020).

two and the precise conditions for when the property is or is not instantiated. Much as I can count as valuing X and believing Y by virtue of my behavior and dispositions, without ever consciously considering the relevant propositions (cf., Audi 1994, Schwitzgebel 2002, Zagzebski 2003), I can count as consenting to an authority by virtue of behavior and dispositions that express the relevant evaluative and cognitive attitudes.[25]

I now want to draw on Greene's account of the legitimacy of governments and produce a general account of legitimacy for norms – or, as I will prefer to put it, an account of their normative authority. Take *quality consent* to mean the kind of consent to be under the authority of a certain norm that comes from recognizing its success in fulfilling its essential function (i.e. consent that is based on telic recognition). Take a norm to have *normative authority* when it is appropriate to criticize, resent, or sanction someone for flouting them. I propose the following principle connecting *consent and authority*:

> **Consent and Authority (C&A):** Norm N has some degree of normative authority over S at t if (a) S is a member of community C, and (b) a sufficiently large number of the members of C have quality-consented to be under the authority of N.

Notice that C&A only lays down a sufficient condition for when it is appropriate to criticize, resent, or sanction someone for flouting a certain norm. I am here only identifying one way in which a norm can come to have normative authority over us.[26] Also notice how C&A identifies individuals as the *source* of normative authority but identifies the communities composed of these individuals as its *grounds*. According to C&A, my quality-consenting to be under the authority of some norm is not enough for it to have normative authority over me, just as my quality-dissenting from some norm is not enough to remove me from under its rule. On this point, C&A contrasts with other claims about authority that take individuals as *both* source and ground, for example:

> The authority of another person is justified for me by my conscientious judgment that if I do what the authority directs (or believe what the authority tells me), the result will survive my own conscientious self-reflection better than if I try to figure out what to do/believe myself. (Zagzeski 2012, 148)

---

[25] This is not, notice, to make quality consent merely hypothetical. The reliance on dispositions is not the same as the claim "he *would* consent, if only…", but is rather the claim "he *does* consent, because…". There is an enormous amount of complexity behind the idea of non-explicit consent. I ignore that complexity for the purposes of the sketch I will offer below.

[26] In particular, C&A does not purport to identify a mechanism by which *normative reasons* are generated. I myself deny that there is a necessary connection between the normative authority of N over S (and the consequent propriety of criticizing, resenting, or sanctioning S for flouting N) and S's normative reasons for Φ-ing. Whether I am right or wrong about this, the point is that this matter is substantive and independent of C&A.

C&A is more radically social than this. It claims instead that the normative authority of some norm can be derived from a similar source but without really depending on it. In other words, C&A harnesses the power that we have as individuals to sometimes imbue normative force to previously neutral rules or policies, but it filters that power through the social structure that a collection of such powers can create.[27]

Finally, notice how C&A takes normative authority to come in degrees and how it leaves unspecified what is the precise function that goes from some ratio of quality-consent within C to some degree of normative authority for N over the members of C. This is by design.

With C&A in hand, I will propose below an account of doxastic justification that is tied to what I will call "epistemic consent" and that shares in the three attractive features of Greene's account mentioned above. By connecting the authority of epistemic norms to the authority we have over ourselves, the account explains why it can be appropriate to criticize, resent, or sanction someone for having unjustified beliefs and it avoids, in this way, both the phenomenon of normative alienation and the phenomenon of normative parochialism. All of this, of course, while respecting the social and collective nature of epistemic justification and thereby avoiding the slide into an implausible kind of subjective normative relativism.

4. Epistemic Consent and Doxastic Justification

Recall the *Normative Authority Conception* of justification I identified at the beginning of this paper as my guide:

> **NAC:** S is justified in their belief that p at t (to some degree n) if and only if their believing that p at t is not ruled out by epistemic norms that have normative authority over S at t.

Differently from traditional conceptions, NAC prioritizes the way in which the state of being justified grounds the appropriateness of our practices of epistemic accountability: proper criticism, resentment, and sanctioning *on the basis of someone's flouting an epistemic norm.* To be doxastically justified in believing that p, in other words, is for it to be *epistemically inappropriate* to criticize, resent, and sanction me (though not necessarily inappropriate on other grounds).

---

[27] For more on this normative power of individuals, see, e.g., Zagzebski (2005, 2012) and Chang (2013a, 2013b).

But NAC is not yet an account of justification. For that, we first need an account of normative authority – an account of what makes the violation of a norm the proper ground for criticism, resentment, and sanctions. In the previous section, however, we have identified a source of normative authority that is a candidate for this post:

> **Consent and Authority (C&A):** Norm N has some degree of normative authority over S at t if (a) S is a member of community C, and (b) a sufficiently large number of the members of C have quality-consented to be under the authority of N.

My suggestion now is that C&A explains the kind of normative authority that is had by the epistemic norms of justification. I call the resulting account the *Epistemic Consent Account* of justification:

> **ECA:** S is justified in their belief that p at t (to some degree n) if and only if, for any epistemic community E that S is a member of, there is no norm N such that (i) a sufficiently large number of the members of E have epistemically quality-consented to be under the authority of N, and (ii) S's believing that p at t is ruled out by N.[28]

According to ECA, facts about doxastic justification are social facts: they are facts about a relation that holds between individuals and their communities. This kind of social-constructivism about epistemic normativity is not new.[29] What is unique to ECA, however, is the claim that the ultimate source of the normative authority of epistemic norms is *epistemic quality-consent*: consent to be under the authority of epistemic norms that comes from recognizing their success in fulfilling their essential function (i.e., epistemic consent that is based on telic recognition). If epistemic norm N is to have normative authority over anyone, according to ECA, then a significant number of individuals must have endorsed N on account of recognizing that N's rule achieves what it essentially purports to achieve. I conclude by saying something about this central and distinctive element.

I take a broad view of the essential function of an epistemic norm – similar to the broad view Greene takes of the essential function of governance – namely, delivering truth-related benefits. Just as an individual's consent-producing endorsement of governance has a fixed target with a flexible content, an individual's consent-producing endorsement of epistemic norms can be based on a recognition that it delivers a variety of epistemic goods. Some may endorse norm N on account of recognizing its

---

[28] ECA is an account of "doxastic justification." As such, it purports to describe the conditions under which someone's state of believing that p is epistemically justified. Epistemologists, however, have often found useful to distinguish between *being justified in believing that p* and *having justification for believing that p*. To give an account of this later state is to give an account of "propositional justification." On some views, propositional justification is the more fundamental notion, with doxastic justification defined partly in terms of it. My view reverses the order of explanation (cf. Goldman 1979). On my view, propositional justification is a relation between S and a proposition, at a time, that is defined in terms of the epistemic status S's belief in p *would have*, at that time, given ECA. See Kornblith's (2022) contribution to this volume for a defense of this reversal and see Silva and Oliveira (forthcoming) for a discussion of these and related issues.
[29] See, e.g., Brandon (1994) and Goldberg (2018).

increase of our access to truth, others for its production of knowledge, and others for its connection to understanding; still others may endorse it out of a concern for coherence, reflective endorsement, evidential support, reliability, intellectual virtue, and so on. In each case, an individual is recognizing that N's rule delivers truth-related benefits and, by endorsing it on these different but related grounds, instantiates epistemic quality consent to be under its authority. The same individual, of course, may also endorse different norms for different truth-related benefits. There is no presumption here that each of us has one and only fundamental epistemic commitment and concern.

Telic recognition, however, is factive. Epistemic quality-consent – consent to be under the authority of an epistemic norm – only occurs when the relevant norm's rule in fact delivers some truth-related benefit, aside from being recognized as such by some individual. This is why it is important to resist the temptation to reduce the essential function of epistemic norms to the delivery of one or another of the epistemic goods we are variously concerned with. Norm N's rule can succeed in delivering truth-related benefits (and be correctly recognized as such) without delivering all of them, and without delivering any particular one of them as well.[30] What makes these benefits "truth-related," in other words, is not their objective relation to a more "fundamental" good of truth – indeed, the fact of the matter seems to be that many of these goods are not so related anyway (cf., e.g., Kornblith 2012) – but rather that our non-instrumental commitment and concern for them is *motivated* by a commitment to, and a concern for, truth.[31] In this way, a religious leader, or a philosophical mentor, may be factually misguided in their teachings and yet deliver the truth-related benefits of coherence, reflective endorsement, evidential support, intellectual virtue, and even understanding, such that the endorsement of an epistemic norm N respecting their authority can be based on the proper kind of telic recognition that grounds epistemic quality-consent, despite N' rule failing to deliver truth.[32]

Importantly – indeed, most importantly – it is precisely by *engaging* in practices of epistemic accountability – by holding ourselves and others accountable to certain epistemic norms – that we instantiate our epistemic consent to be under the normative authority of those norms. Just as ownership or commitment to reasons, and just as political consent, epistemically consenting to be

---

[30] This suggestion is in tension with a common way of thinking about epistemic goods, a way that is well represented by Pritchard (2014, 113) when he says that "from a purely epistemic point of view it is ultimately only truth that we should care about."

[31] My commitment and concern for coherence or reflective endorsement, for example, is non-instrumentally truth-related, in this sense, when (a) I do not value coherence or reflective endorsement *as a means to* the further end of truth, and yet (b) the motivational structure that causes such a non-instrumental concern is itself grounded in a prior commitment to, and concern for, truth.

[32] For arguments suggesting that understanding is possible with and through falsehoods (and, indeed, sometimes *only* through them), see Zagzebski (2001) and Elgin (2017). For an argument that the kind of reflective equilibrium we find in Hume, Popper, Goodman, and Quine disconnects justification from likelihood of truth, see Johnsen (2017).

under the authority of some epistemic norm does not require explicit and overt affirmative expressions – or similarly structured thoughts. I can count as having the required cognitive and evaluative states behind consent merely by actually holding myself and others accountable to an epistemic norm, provided I do so out of a commitment to, and concern for, the truth-related benefits I endorse.[33] The idea here is that when I sense that something is *epistemically* off and feel the pull to do something about it, I am thereby revealing that by my own epistemic lights – lights that come from my commitments to, and concerns for, truth – I seem to be in violation of norms that deliver truth-related benefits (broadly construed). Similarly, when I hold you accountable for being *epistemically* off, I am expressing my judgment that you have flouted a norm that delivers truth-related benefits as well – a judgment that arises out of my commitment to, and concern for, truth. Here, notice, there is no reason to worry about the imperfect match between overt behavior (be it internal or external) and quality consent. The behavior we are considering here is not my *mere obedience or compliance* to the relevant norms (as if by accident or by virtue of some other concern), but rather my *motivated obedience* and *voluntary enforcement* of those epistemic norms on myself and others.

We finally have in view the full nature of the explanatory reversal I mentioned at the beginning of the paper. According to NAC, to be doxastically justified in my believing that p (to some degree n) is for it to be *epistemically inappropriate* to criticize, resent, or sanction me. There are, therefore, facts about the propriety of holding me epistemically accountable in these ways – facts that determine whether I am justified in my belief. The surprising claim we get from combining NAC with C&A, however, is that what grounds and explains these facts about the propriety of epistemic accountability are simply further facts about our practices of epistemic accountability themselves. For it to be epistemically inappropriate to hold me accountable for believing that p at t, after all, is just for me to be part of an epistemic community that has *epistemically quality-consented* to be under epistemic norms that do not rule out my believing that p at t. And, on the view I am urging here, an epistemic community epistemically quality-consents to be under certain epistemic norms simply by expressing their fundamental commitment to, and concern for, truth in practices of epistemic accountability in the first place. For better or for worse, this is truly epistemic normativity from the ground up.[34]

---

[33] There is, of course, much in our practices of accountability that are not based on a concern for truth-related benefits. These do not count, on my terminology, as practices of *epistemic* accountability, even when their targets are doxastic states. How to properly characterize the relevant commitment to, and concern for, truth that gives rise to practices that are non-instrumentally attuned to truth-related benefits – in order to distinguish those that are practices of epistemic accountability from those that are not – is a complex issue that I also leave open another day. See, however, Zagzebski (2003, 178-81) for a discussion of "epistemic motivation," and Zagzebski (2005, 186-192) for a discussion of the "logic of caring" for epistemic goods.

[34] In similar but different ways, both Strawson and Alston take the relevant practices in their purview to be explanatorily prior to the grounds of their propriety. For Strawson, as Beglin (2018, 618) tells us, "our responsibility practices are an expression of certain concerns, and they thus reflect these concerns and are answerable to them." For Alston (1991, 158), "what is, factually, a more or less fixed habit of going from inputs of type I to a belief output of correlated type B, is also,

Of course, a satisfying defense of ECA requires saying much more than this about consent and other matters. How do we individuate epistemic communities?[35] How do we extract "the norms that I am holding myself and others accountable to" from the complex patterns of internal and external behavior that constitute our practices of epistemic accountability?[36] What are the norms, after all, that have normative authority over me and most of us?[37] I do not have the space to address these and other questions here. But there are clearly different possible answers. I hope this paper is just the beginning of a conversation on how to give ECA its most plausible articulation.

Conclusion

If we take doxastic justification as the inchoate ideal behind our practices of epistemic accountability, then we are faced with the problem of explaining how the norms of justification have the kind of normative authority over us that makes it appropriate to criticize, resent, and sanction someone for flouting them. I have argued here that we can make progress on this problem by paying close attention to work in political philosophy connecting the legitimacy of governments to consent to be ruled. The resulting account of doxastic justification – ECA – has four features that I find attractive in an account of doxastic justification.

First, and most obviously, ECA produces an account of justification that is connected to an account of normative authority that I find plausible. It explains, to my liking, why it is appropriate to criticize, resent, and sanction me when I have an unjustified belief: I belong to an epistemic community, I contributed to the grounds of the normative authority of our shared epistemic standards (by having truth-related commitments and concerns and by expressing them through various internal and

---

evaluatively, a principle of justification for beliefs so formed. The principle says that when a belief of type B is formed on the basis of an input of type I, that belief is thereby (prima facie) justified." My view owes much to theirs. Goldberg (2018, 147-9) endorses a similar reversal, explaining epistemic standards by way of the epistemic expectations we are entitled to have of each other. By constraining the normative power of our actual practices with the normative notion of "entitlement," however, Goldberg's view is not, normatively speaking, "practices-first" all the way down.

[35] In order to avoid a worry similar to the one I raised for constitutivists in section 3, answering this question requires attention to how membership into an epistemic community is itself an expression of an individual's fundamental commitment to, and concern for, truth. In particular, it requires attention to the difference between epistemic dissent and epistemic defection.

[36] Answering this question requires the careful work of understanding how our fundamental commitment to, and concern for, truth is *actually expressed* in our various practices of epistemic accountability. And while for the moment I stand with Strawson (1962, 2) – "I can give no simple description of the field of phenomena at the centre… for the field is too complex" – a good start is Brown's (2019, 2020) work on epistemic blame and Goldberg's (2018, ch. 5) work on epistemic social expectations.

[37] ECA represents a kind of naturalistic approach to epistemology. There is no telling what the norms of justification are *in advance*, as it were, of doing the careful and empirical work of examining and understanding our practices of epistemic accountability. That said, my sense is that this approach reveals that all norms of justification are what Josh DiPaolo (2019, 2049) calls "norms of compensation." See the second feature of ECA that I highlight in the conclusion below.

external practices of accountability), and I have failed them. This is the sense in which ECA avoids the phenomena of normative alienation.[38]

Second, ECA achieves this without imposing the normative authority of certain norms as "universal," "self-evident," or "brute" facts. The similarities between our natures, our fundamental concerns, and our circumstances, no doubt ensures a great amount of overlap between the kinds of epistemic norms that have normative authority over the members of different epistemic communities – even at different times and places. But finer differences in our cultural upbringings, or finer differences in our concerns and contexts, may well produce some amount of divergence on their details. This is not a matter of some communities endorsing implausible anti-inductive norms or some such. (Would a group's commitment to, and concern for, truth really ever be expressed in practices of accountability that are guided by anti-inductive norms?) Instead, this is a matter of different communities having different evidential standards and thresholds, different views on the proper place for trust on epistemic authorities, different practices regarding disagreement, different habits for tolerating uncertainty, inconsistency, groundlessness, and so on. Much like the non-subjectivist relativism of Greene's account of political legitimacy, we here have the appropriate kind of normative flexibility: understanding what I am properly accountable for, epistemically speaking, requires first understanding the norms of the epistemic communities I am embedded within. This is the sense in which ECA avoids the phenomena of normative parochialism.[39]

Third, ECA delivers a view of doxastic justification – and a related understanding of our practices of epistemic accountability – that respects the fact that we lack sufficiently meaningful doxastic agency. There is no reason to think that we cannot consent to be under the authority of norms that govern states of ours for which we have minimal amounts of agency. Those of us, like myself, who think that our degree of doxastic agency is closer on the spectrum to our degree of digestive agency than to the degree of agency we typically have over our bodily movements should welcome this result.[40]

---

[38] *Subjective permissivists* about evidential support also suggest ways in which individuals can play a role in determining the details of their epistemic situations (see, e.g., Schoenfield 2014, Callahan 2019). Their views, however, typically take for granted the normative authority of evidentialist norms and only give the subject enough powers to affect the particular prescriptions that apply to them. ECA, on the other hand, takes the normative authority of no epistemic norm for granted.
[39] The worry about parochialism is not political. What would be the grounds for excluding some set of epistemic norms *a priori* as unacceptable if not the privileging of another set of norms? And what non-arbitrary, non-infinitely-regressable, and non-circular grounds could one have for privileging one set of epistemic norms over another? (cf. Alston 1989, 10-11; 1991, 149-153).
[40] In Oliveira (2015), I resisted an argument against the basing requirement on justification (found in Silva 2015), by drawing an agential contrast between the epistemic and moral domain and by articulating a notion of "non-agential permissibility" that made the norms of epistemic evaluation more like the norms for evaluating clocks than the norms for evaluating actions. I take the present paper to be a development of that central notion and a correction of some claims I made in its defense. In particular, I was wrong in thinking that non-agential permissibility made talk of praise and blame inappropriate, and incorrect in thinking that doxastic justification necessarily has basing requirement. See Silva's (2022)

Fourth, and this is perhaps the most controversial of these attractive features, ECA allows for the possibility of "unclear" and "borderline" cases. These kinds of cases may arise in at least two kinds of circumstances: when it is unclear whether the number of members of a certain epistemic community who have epistemically consented to N is sufficiently large, and when one is a member of two or more communities with conflicting norms. Yet I think unclarity about whether I am doxastically justified in certain circumstances is precisely the right result. If there is a case for imprecision regarding the very concept of belief (cf. Schwitzgebel, 2002, 252), there is certainly a strong case for imprecision in our concept of justification. In my view, we shouldn't expect exactness in matters of normative evaluation – the kind of exactness that tends to come, for example, from modeling the norms of justification on mathematical or scientific (probabilistic) reasoning. Caution and modesty in our attributions of justification should not only be recommended on account of our poor access to the relevant facts, it should also be recommended on account of the nature of the facts we purport to make claims about.[41]

References

Alston, William (1985). "Concepts of Epistemic Justification," *The Monist* 68(1): 57-89.

Alston, William (1989). "A "Doxastic Practice" Approach to Epistemology," in *Knowledge and Skepticism*, edited by Marjorie Clay and Keith Lehrer. Routledge: 1-29.

Alston, William (1991). *Perceiving God: The Epistemology of Religious Experience*. Cornell University Press.

Alston, William (2005). *Beyond "Justification": Dimensions of Epistemic Evaluation*. Cornell University Press.

Armstrong, David (1973). *Belief, Truth, and Knowledge*. Cambridge University Press.

Audi, Robert (1993). *The Structure of Justification*. Cambridge University Press.

Audi, Robert (1994) "Dispositional Beliefs and Dispositions to Believe," *Nous* 28, pp. 419–434.

Audi, Robert (2020). *Seeing, Knowing, and Doing: A Perceptualist Account*. Oxford University Press.

Baxter, Donald (2018). "A Pyrrhonian Interpretation of Hume on Assent," in *Skepticism: From Antiquity to the Present*, edited by Diego Machuca and Baron Reed. Bloomsbury Academic: 380-394.

Beglin, David (2018). "Responsibility, Libertarians, and the "Facts as We Know Them": A Concern-Based Construal of Strawson's Reversal," *Ethics* 128: 612-625.

Bennett, Jonathan (1990). "Why is Belief Involuntary?" *Analysis* 50: 87-107.

Brandom, Robert (1994). *Making it Explicit: Reason, Representing, and Discursive Commitment*. Harvard University Press.

contribution to this volume for discussion. In Oliveira (2021), however, I argued that our lack of sufficiently meaningful doxastic agency means there are no such things as doxastic obligations or even normative reasons to believe. The account of epistemic normativity developed here is intended as complimentary to those results.

[41] This does not mean that ECA entails relativism about attributions of justification – where whether S is justified in their belief that p depends on which community we explicitly or contextually select as providing the standards for evaluation. See Goldberg (2018, 239-243) for a discussion of this issue.

Breyer, Daniel (2010). "Reflective Luck and Belief Ownership". *Acta Analytica*, 25: 133-154.

Breyer, Daniel and Greco, John (2008). "Cognitive Integration and the Ownership of Belief: Response to Bernecker." *Philosophy and Phenomenological Research* 76 (1):173-184.

Brown, Jessica (2019). "Epistemically Blameworthy Belief," *Philosophical Studies* 177: 3595–3614.

Brown, Jessica (2020). "What is Epistemic Blame?" *Nous* 54(2): 389-407.

Callahan, Laura (2019). "Epistemic Existentialism," *Episteme*: 1-16.

Chang, Ruth (2013a) "Grounding Practical Normativity: Going Hybrid," *Philosophical Studies* 164(1): 163–87.

Chang, Ruth (2013b) "Commitments, Reasons, and the Will," in *Oxford Studies in Metaethics (Vol. 8)*, edited by Russ Shafer-Landau. Oxford University Press: 74-113.

Cote-Bouchard, Charles (2015). "Epistemic Instrumentalism and the Too Few Reasons Objection," *International Journal of Philosophical Studies* 23(3): 337-355.

DiPaolo, Joshua (2019). "Second Best Epistemology: Fallibility and Normativity," *Philosophical Studies,* 176: 2043-2066.

Duke-Yonge, Jennifer. (2012). "Ownership, Authorship and Externalist Justification". *Acta Analytica*, 28: 237-252.

Elgin, Catherine (2017). *True Enough*. MIT Press.

Enoch, David (2006). "Agency, Shmagency: Why Normativity Won't Come from What is Constitutive of Action," *Philosophical Review* 115(2): 169–198.

Enoch, David (2011). "Shmagency Revisited". In *New Waves in Metaethics*, edited by Michael Brady. Palgrave Macmillan: 208–233.

Feldman, Richard and Conee, Earl (1985/2004). "Evidentialism," reprinted in *Evidentialism: Essays in Epistemology*. Oxford University Press: 83-100.

Fischer, John and Ravizza, Mark (1998). *Responsibility and Control: A Theory of Moral Responsibility*. Cambridge University Press.

Flowerree, A.K. (2018). "Epistemic Shmagency" In *Metaepistemology: Realism and Anti-Realism*, edited by. Christos Kyriacou and Robin McKenna. Palgrave Macmilllan: 289-310.

Frankish, Keith (2007). "Deciding to Belief Again," *Mind* 116(463): 523-548.

Goldberg, Sanford (2018). *To the Best of Our Knowledge*. Oxford University Press.

Goldman, Alvin (1979/1992). "What is Justified Belief?" Reprinted in *Liaisons: Philosophy Meets the Cognitive and Social Sciences*. MIT Press: 105-126.

Greene, Amanda R. (2016). "Consent and Political Legitimacy," in *Oxford Studies in Political Philosophy*, 2: 71-97.

Greene, Amanda R. (2017). "Legitimacy without Liberalism: A Defense of Max Weber's Standard of Political Legitimacy". *Analyse & Kritik* 39 (2): 295–323.

Greene, Amanda R. (2019). "Is Political Legitimacy Worth Promoting?" *NOMOS: Journal of the American Society for Political and Legal Philosophy* LXI: 65-101.

Greene, Amanda R. (2020). "When Are Markets Illegitimate?" *Social Philosophy and Policy* 36(2): 212-241.

Grimm, Stephen (2009). "Epistemic Normativity," in *Epistemic Value*, edited by Adrian Haddock, Alan Millar, and Duncan Pritchard. Oxford University Press: 243-264.

Hieronymi, Pamela (2008), "Responsibility for Believing," *Synthese* 161: 357-373.

Huemer, Michael (2012). *The Problem of Political Authority*. Palgrave Macmillan.

Johnsen, Bredo (2017). *Righting Epistemology: Hume's Revolution*. Oxford University Press.

Kornblith, Hilary (1993). "Epistemic Normativity," *Synthese* 94(3): 357-376.

Kornblith, Hilary (2012). *On Reflection*. Oxford University Press.

Kornblith, Hilary (2022). "What does Logic Have to do With Justified Belief? Why Doxastic Justification is Fundamental," in *Propositional and Doxastic Justification: New Essays on Their Nature and Significance*, edited by Luis Oliveira and Paul Silva. Routledge.

Korsgaard, Christine (2008). *The Constitution of Agency: Essays on Practical Reason and Moral Psychology*. Oxford: Oxford University Press.

Leite, Adam (2004). "On Justifying and Being Justified," *Philosophical Issues*, 14: 219-253

McCormick, Miriam S. (2011). "Taking Control of Belief," *Philosophical Explorations*, 14 (2): 169-183

McCormick, Miriam S. (2015). *Believing Against the Evidence: Agency and the Ethics of Belief*. Routledge.

McHugh, Conor and Whiting, Daniel (2014). "The Normativity of Belief," *Analysis*, 74(4): 698-713.

Nolfi, Kate (2015) "How to be a Normativist About Belief," *Pacific Philosophical Quarterly,* 96(2): 181-204.

Nozick, Robert (1981). *Philosophical Explanations*. Harvard University Press.

Oliveira, Luis R.G. (2015) "Non-Agential Permissibility in Epistemology," *Australasian Journal of Philosophy* 93 (2): 389-394.

Oliveira, Luis R.G. (2017). "Deontological Evidentialism, Wide-Scope, and Privileged Values," *Philosophical Studies* 174 (2): 485-506.

Oliveira, Luis R.G. (2018) "Ampliative Transmission and Deontological Internalism," *Pacific Philosophical Quarterly*. 99 (2): 174-185.

Oliveira, Luis R.G. (2021) "Evading the Doxastic Puzzle by Deflating Epistemic Normativity," in *Epistemic Duties: New Arguments, New Angles*, edited by Scott Stapleford and Kevin McCain. Routledge: 44-62.

Osborne, Robert (2021) "Doxastic Responsibility, Guidance Control, and Ownership of Belief" Episteme 18(1): 82-98.

Owen, David (2003). "Locke and Hume on Belief, Judgment, and Assent," *Topoi* 22(1): 15-28.

Plantinga, Alvin (1990). "Justification in the 20th Century," *Philosophy and Phenomenological Research* 50: 45-71.

Perrine, Timothy (2018). "Justification, Justifying, and Leite's Localism," *Acta Analytica*, 33 (4): 505-524.

Perrine, Timothy (forthcoming). "Conceptions of Epistemic Value," *Episteme*.

Pritchard, Duncan (2005). *Epistemic Luck*. Oxford University Press.

Pritchard, Duncan. (2014). "Truth as the Fundamental Epistemic Good." In *The Ethics of Belief*, edited by Jonathan Matheson and Rico Vitz. Oxford University Press: 112-129.

Rawls, John (1996). *Political Liberalism*. Columbia University Press.

Raz, Joseph (1988). *The Morality of Freedom*. Oxford University Press.

Scanlon, T. M. (1998). *What We Owe Each Other*. Harvard University Press.

Schoenfield, Miriam (2014) 'Permission to Believe: Why Permissivism is True and What it Tells us About Irrelevant Influences on Belief.' *Nous* 48(2), 193–218.

Scott-Kakures, Dion (1994). "On Belief and the Captivity of the Will," *Philosophy and Phenomenological Research*, 54: 77-103.

Shah, Nishi, and David Velleman. (2005). "Doxastic Deliberation". *Philosophical Review* 114 (4): 497–534.

Silva, Paul Jr. (2015). "Does Doxastic Justification Have a Basing Requirement?" *Australasian Journal of Philosophy*, 93(2): 371-387.

Silva, Paul Jr. (2022). "Does the Basing Demand on Doxastic Justification Have Any Dialectical Force? A Response to Oliveira," in *Propositional and Doxastic Justification: New Essays on Their Nature and Significance*, edited by Luis Oliveira and Paul Silva. Routledge.

Silva, Paul Jr. and Oliveira, Luis R.G. (forthcoming). "Propositional Justification and Doxastic Justification," in *Routledge Handbook of the Philosophy of Evidence*, edited by M. Lasonen-Aarnio and C. Littlejohn. Routledge.

Siscoe, Robert Weston (2021). "Belief, Rational and Justified," *Mind*. 130 (517): 59-83.

Smith, Angela (2012). "Attributability, Answerability, and Accountability: In Defense of a Unified Account." *Ethics*, 122(3): 575-589.

Sosa, Ernest (2007). *A Virtue Epistemology*. Oxford University Press.

Strawson, P. F. (1962). "Freedom and Resentment," *Proceedings of the British Academy* 48: 1-25

Street, Sharon (2009). "Evolution and the Normativity of Epistemic Reasons," *Canadian Journal of Philosophy*, Vol. 39, supplement 1: 213-248.

Schwitzgebel, Eric (2002) "A Phenomenal, Dispositional Account of Belief," *Nous* 36(2): 249–275.

Todd, Patrick (2016). "Strawson, Moral Responsibility, and the 'Order of Explanation': An Intervention," *Ethics* 127: 208-40.

Velleman, David. 2009. *How We Get Along*. Cambridge University Press.

Watson, Gary. (2004). *Agency and Answerability*. Oxford: Oxford University Press.

Williams, Bernard (1973). "Deciding to Believe," in *Problems of the Self*. Cambridge University Press.

Williamson, Timothy (2007). *The Philosophy of Philosophy*. Wiley-Blackwell.

Winters, Barbara (1979). "Believing at Will," *The Journal of Philosophy*, 76(5): 243-256.

Zagzebski, Linda (1996). *Virtues of the Mind*. Cambridge University Press.

Zagzebski, Linda (2001/2020) "Recovering Understanding," reprinted in *Epistemic Values*. Oxford University Press: 57-77.

Zagzebski, Linda (2003/2020). "Intellectual Motivation and the Good of Truth," reprinted in *Epistemic Values*. Oxford University Press: 168-185.

Zagzebski, Linda (2005/2020). "Epistemic Value and the Primacy of What We Care About," reprinted in *Epistemic Values*. Oxford University Press: 186-203.

Zagzebski, Linda (2012). *Epistemic Authority*. Oxford University Press.

Zagzebski, Linda (2020). *Epistemic Values: Collected Papers in Epistemology*. Oxford University Press.