

On the Relevance of Self-Disclosure for Epistemic Responsibility

Daniel Buckley, Penn State Harrisburg

Forthcoming in *Ergo* (please cite the published version)

Abstract: A number of authors have argued that, in order for S to be appropriately held morally responsible for some action or attitude (say, via moral blame), that action or attitude must somehow reflect or express a negative aspect of S's ("true", "deep", or "real") self. Recently, theorists of "epistemic blame" and "epistemic accountability" have also incorporated certain "self-disclosure" conditions into their accounts of these phenomena. In this paper, I will argue that accounts of epistemic responsibility which require disclosure of an objectionable feature of a person's self neglect central aspects of our responsibility practices. Specifically, I will argue that we can appropriately hold individuals epistemically responsible for things which do not reflect or express negative aspects of "who they are" in any way. While disclosure of an objectionable feature of (what I will call) a person's "doxastic self" isn't altogether irrelevant when it comes to holding that person epistemically responsible, it also isn't necessary. One lesson I wish to draw is that, insofar as our epistemic responsibility practices are capable of shedding light on the nature of epistemic normativity itself, theorists interested in the latter should be focused on more than just accountability responses which track objectionable features of a person's self.

Keywords: Epistemic Responsibility, Self-Disclosure, Epistemic Normativity, Epistemic Norms, Agency and Responsibility

1. Responsibility and Self-Disclosure

What is it for a subject S *to be* morally responsible for some action or attitude, as distinct from S's being *held* morally responsible for that action or attitude by other people? The former, one might think, is a necessary condition for the appropriateness of the latter. It seems that a person can't be appropriately *held* morally responsible (say, through blame or reproach) unless what they're being held responsible *for* bears some kind of appropriate connection to their agency.

Consider two different vase smashers. The first vase smasher trips while walking and accidentally knocks over an expensive vase. The second vase smasher picks up the vase and intentionally smashes it on the ground out of malice. It seems that we shouldn't hold the first vase smasher morally responsible; it wouldn't be appropriate to blame or criticize this person, morally speaking (assume that the person wasn't clumsy, inattentive, drunk, etc.). Moreover, it seems that part of the reason *why* we shouldn't do so is because the action isn't connected to this person's agency in the right kind of way.

But what *is* this "appropriate connection" that must obtain between our agency and the things that we do and think in order for moral accountability responses to be on the table? According to "self-disclosure" theorists, in order for S to be an appropriate target of moral accountability responses, what she

is being held accountable *for* must somehow express or disclose S's ("true", "deep", or "real") self.¹ When the second vase smasher smashes the vase, this action tells us something about *who the person is*; it is disclosive of their adopted ends, and it expresses their evaluative commitments and overarching orientation. While these ends, values, and commitments are clearly morally *objectionable* in the case at hand, they are still *his*. According to self-disclosure theorists, this is the kind of connection that must obtain between our agency and the things that we do and think in order for moral accountability responses like blame and reproach to be on the table.² By contrast, when the first vase smasher smashes the vase, this tells us nothing about the person; it does not reveal or express "who they are" in any significant sense. For all we know, this person could have impeccable ends, values, and commitments, morally speaking.

Self-disclosure views are often put forward as compatibilist-friendly approaches to the issue of moral responsibility. Proponents of the view hold that there are certain things which are "attributable" to us as agents, in a sense of this word that runs deeper than mere true predication (e.g. "Paul has big ears"), whether or not the thesis of determinism is true. As Watson (2004) has noted,

The point of speaking of the "real self" is not metaphysical, to penetrate to one's ontological center; what is in question is an individual's fundamental evaluative orientation. (Watson 2004: 271)

Because the "real self" implicates one's identity and evaluative orientation, Watson argues that it has "ethical depth in an obvious sense" (2004: 271). Things which are "attributable" to us, in this richer sense of the word, might include various states of mind, character traits, as well as actions which stem from these. As Shoemaker (2015) has noted, the class of the "attributable" consists in

¹ According to John Dewey, "We are responsible for our conduct because that conduct is ourselves objectified in actions." (Dewey 1891:161). Peter Strawson could also be considered an adherent of the view insofar as he saw the reactive attitudes (e.g. resentment and indignation) as tracking the "qualify of one's will" (Strawson 1962/2003). Other early proponents include: Frankfurt (1971), Taylor (1976) and Watson (1975). More recent defenses include: Scanlon (1998), Shoemaker (2015), Smith (2005; 2015), and Sripada (2016). See: Reis-Dennis (2018) for a recent criticism of the view.

² This is often taken to be a necessary but not sufficient condition. One could hold that further conditions must be met in order for certain forms of accountability (e.g. blame, sanction, etc.) to be called for. Coates (2020) makes this point by distinguishing between the fittingness of an accountability response, on the one hand, and its all-things-considered appropriateness, on the other. Considerations which might be factored in when determining the all-things-considered appropriateness of an accountability response include: "standing" considerations, i.e. considerations pertaining to the kind of relationship that obtains between blamer and blamee; fairness considerations, e.g. whether or not one had a reasonable opportunity to avoid incurring a sanction; and communicative considerations, e.g. whether the target of blame has the capacities and abilities required for understanding the moral demands implicit in various forms of protest, reproach, or criticism.

a privileged subset of our psyche – our representative voice for purposes of responsibility – which includes some attitudes but excludes many others. The excluded elements are among the following: (a) pains, itches, and reflexes; (b) what I call psychic junk, e.g., the obsessional, compulsive, and addictive urges and cravings of many psychologically disordered people, as well as the random thoughts, impulses, and images that cross the consciousness of everyone from time to time; and (c) some attitudes formed on whims or as the result of altered states of consciousness. (Shoemaker 2015: 118)

Unsurprisingly, the view comes in many different varieties. In particular, self-disclosure theorists are divided when it comes to what it takes for something *to be* disclosive or expressive of one’s “true” or “real” self.³ I will have to set these important disagreements to the side in order to convey the basic thrust of the view. What proponents of the view hold is that attitudes and actions are connected to our agency in the appropriate way – and thus meet the basic condition of moral responsibility – when they somehow disclose or express our adopted ends and evaluative commitments, i.e. when they somehow reveal “who we are”.

In what follows, I will be interested in self-disclosure as it relates specifically to *epistemic* responsibility. Just as we can express or reveal “who we are” in action, it seems that we can also express “who we are” in thought or belief. In other words, it appears as though our doxastic attitudes (e.g. belief, disbelief, and suspension of judgment), as well as acts associated with the regulation of such attitudes, are “attributable” to us in the sense required for self-disclosure. A person’s beliefs aren’t like pains, itches, or mere “psychic junk”; they comprise part of an individual’s take on the world, they are responsive to our evaluative judgments and assessments of reasons, and they are indicative of general traits of character and various casts of mind.⁴

Consider some observations in connection with this last point: An individual might be *hasty* or *lazy* in believing a certain claim; someone might be *dogmatic* or *biased* in maintaining a certain belief; one might be *diligent*, *careful*, or *charitable* when managing their beliefs; and so on. I take the italicized words to be apt ways of characterizing an individual qua believer. In other words, these are ways of

³ Scanlon (1998), Smith (2005; 2015), and Watson (1975) defend rationalistic approaches where the “real” self is tied to our evaluative judgments and assessments of reasons; Frankfurt (1971) understands the real self in terms of what he calls “second-order volitions”, e.g. a desire that some first-order desire be one’s will; Sripada (2016) equates the real self with one’s “cares”, i.e. “a complex syndrome of motivational, commitmental, evaluative, and affective dispositions” (2016: 1211); and Shoemaker (2015) defends an ecumenical approach where the real self can be expressed either in one’s cares or in one’s evaluative judgments.

⁴ See: Hieronymi (2008), McHugh (2013), and Smith (2005) for defense of the view that our doxastic attitudes are “attributable” to us in the relevant sense.

characterizing what we might call an individual's "doxastic self", i.e. their "self" as it pertains to the attitude of belief.⁵

The question I would like to consider here is the following: For any object of epistemic evaluation X and subject S, is it the case that, in order for S to be held epistemically responsible for X, X must express or disclose some negative aspect of S's doxastic self? I will defend a negative answer to this question: A person can be held epistemically responsible for things which are in no way expressive or disclosive of an objectionable "doxastic self". In other words, I will argue that there are cases where all of (1)-(3) are met:

(1) An object of epistemic evaluation X (e.g. a doxastic attitude) is attributable to a person S.⁶

(2) X does not reveal or express anything negative about S's "doxastic self" (when the latter is assessed from the distinctly epistemic point of view).

(3) S can be appropriately held epistemically responsible for X.

I will assume that "attributability" in the form of self-disclosure, on *some* version of the view, successfully captures the basic condition of responsibility.⁷ This is a controversial assumption, but I do not have the space to defend it here. Instead, I would like to consider whether or not that which is attributable to a person must *also* reveal or express an objectionable "doxastic self" in order for epistemic accountability responses to be on the table. I will argue that this is not so. If epistemic norms underlie our legitimate practices of interpersonal epistemic responsibility, and if I'm right that there are cases where (1)-(3) are all met, then this implies that individuals can be the appropriate targets of accountability responses for epistemic norm violations which in no way reflect poorly on "who they are" qua believers. As we will see, when it comes to cases where (1)-(3) are met, it will not be apt to characterize these accountability responses as "blame". However, I will argue that, even in such cases, these responses constitute a genuine form of accountability nonetheless.

⁵ Note that these ways of characterizing a person's doxastic self can be true even if we can't believe at will, i.e. even if doxastic voluntarism is false.

⁶ I will take doxastic attitudes and acts associated with the regulation of such attitudes (e.g. inference, judgment, reasoning, etc.) to be the paradigmatic objects of epistemic evaluation. However, I do not mean to restrict "epistemic" evaluation to these. For instance, acts of evidence-*gathering* might also be open to epistemic evaluation. See: Flores & Woodard (Forthcoming) and Freidman (Forthcoming) for defenses of this "expansive" view of epistemic evaluation.

⁷ In particular, I'm inclined to think that, when it comes to responsibility for doxastic attitudes, "answerability" views of the kind defended by Scanlon (1998) and Smith (2005; 2015) are quite promising. These views understand the basic condition for responsibility in terms of the in-principle intelligibility of a certain kind of justificatory request, viz. a request for normative reasons.

One clarification concerning (2) before proceeding. I will leave open the possibility that there are different ways in which our “doxastic selves” can be assessed. For instance, our doxastic selves might be evaluable from the *moral* or *practical* points of view. My focus will be on evaluations of our doxastic selves from the *distinctly epistemic* point of view; roughly, the point of view which privileges things like knowledge, truth, and evidence.⁸ I will not here develop a full account of what a good doxastic self looks like from this point of view. I will assume that characteristic markers of such a self will include things like: making a good faith effort to avoid bias in one’s cognitive life, manifesting care and attention vis-à-vis truth-relevant considerations (e.g. evidence) when regulating one’s beliefs, exhibiting proficiency at reasoning, and maintaining a willingness to reflect on, and provide (epistemic) justification for, one’s doxastic attitudes.⁹

My argument will center around a pair of cases, Case 1 and Case 2. Case 1 will involve an individual who, in believing p, reveals an objectionable doxastic self from the epistemic point of view. Case 2 will involve an individual who, in believing p, does not reveal an objectionable doxastic self from the epistemic point of view. I will argue that, in spite of not revealing an objectionable doxastic self, the individual in Case 2 can be held epistemically responsible for believing p (albeit in a way that perhaps differs from Case 1). Thus, I will argue that all of (1)-(3) are met in Case 2. Here are the cases:

Case 1: Tom has an abundance of available evidence which is such that, were he to consider it more carefully, he would feel rationally compelled to abandon his persistent denial in the existence of human-induced climate change. Nevertheless, Tom doesn’t consider this evidence more carefully and, as a result, he maintains his belief that human-induced climate change isn’t real.

Case 2: Bill lives in a community run by a group of individuals who have managed to shut off access to evidence from the outside world pertaining to the existence of human-induced climate change. Bill is interested in forming true beliefs about this topic, frequently going to the library to research it. Nevertheless, given the evidence available to him, Bill believes that human-induced climate change isn’t real.

⁸ This is not to deny that “epistemic” evaluation is tied to the social realm in important ways; I’m inclined to think that it is. It is to say that, even if social, epistemic evaluation of a person is distinguishable from moral assessment.

⁹ While there will be some overlap, I would like to distinguish epistemically good “doxastic selves” from the intellectually virtuous traits of character that are of interest to some virtue theorists. The key difference is that a good “doxastic self” might be manifested in a fleeting, one-off, way (e.g. by someone who is *typically* intellectually vicious). Thus, while an epistemically good “doxastic self” might be a fairly fixed or stable feature of one’s psychology, it need not be. One can, in a moment of uncharacteristic patience, attention, or open-mindedness, *summon* or “put on” an epistemically good doxastic self. Alternatively, a person who is typically intellectually virtuous might manifest an epistemically bad “doxastic self” on a specific occasion.

Once again, in claiming that (1)-(3) are met in Case 2, I do not mean to suggest that Bill is *blameworthy* for believing as he does; I do not think that he is. Rather, I will argue that the epistemic domain involves accountability responses which can be adopted towards an individual even when (as in Case 2) the individual does not reveal or disclose an objectionable doxastic self from the epistemic point of view. When these responses are taken up in such cases, a person is held epistemically responsible without also being blamed.

To be clear, my aim is not to deny the existence of “epistemic blame”. What I say here is compatible with the existence of a distinctly epistemic form of blame. For instance, I’m open to the possibility that Tom in Case 1 is open to epistemic blame. Rather, my aim is to show that there are forms of interpersonal accountability in the epistemic domain which are not correctly classified as “blame” and which also do not require disclosure of an objectionable doxastic self in order to be appropriately adopted. To the extent that theorists of epistemic blame have neglected this aspect of our responsibility practices, their accounts are *incomplete*, not necessarily incorrect.

Before arguing that Case 2 represents an instance where (1)-(3) are all met, I would first like to say more about what it is to “hold” someone responsible for something, and specifically what it is to “hold” someone *epistemically* responsible for something. In the next section (sec. 2) I will turn to this issue. After that, I will consider what kinds of circumstances appropriately give rise to epistemic accountability responses. In sec. 3. I will show that a number of recent accounts of “epistemic blame” and “epistemic accountability” incorporate conditions which require disclosure of an objectionable doxastic self. In sec. 4, I will challenge this consensus by arguing that (1)-(3) all obtain in Case 2.

2. On Being vs. Holding (Epistemically) Responsible

We need to distinguish “holding” a person S responsible for X from S’s *being* responsible for X. The appropriateness conditions of the former are more stringent than the conditions of the latter.¹⁰ Self-disclosure theorists hold that S is responsible for X in the second, less stringent sense, just in case X is “attributable” to S in the manner discussed above, i.e. X somehow reveals or discloses S’s values, ends, and commitments. Meeting this condition is what *opens up the possibility* of S’s being “held” responsible. But what is to “hold” S responsible for X (also known as holding S accountable for X)?

The first thing to note here is that “holding” someone responsible goes beyond mere evaluative assessment.¹¹ Return once again to our malicious vase smasher from earlier. Imagine a detached third-

¹⁰ See note 2.

party observer who evaluates the vase smasher as vicious and cruel, but who does nothing more than this. In particular, say that this observer doesn't also: undergo any reactive emotions (e.g. resentment or indignation), protest the behavior, verbally reproach the vase smasher, demand recompense or apology, change their plans and intentions vis-à-vis the vase smasher going forward, etc. This detached observer still *attributes* the smashing of the vase to the individual; she still sees the act as issuing from the person's values and commitments. Moreover, insofar as she sees these values and commitments as cruel and vicious, she is engaged in a kind of negative evaluative assessment; one which implies responsibility in the minimal, "attributability" sense discussed earlier. Nevertheless, she does not "hold" the vase smasher responsible. The latter requires a certain kind of *engagement* with the offender which the detached observer hasn't yet exhibited (cf. the distinction between judging something to be valuable and actually valuing it; the former can obtain without one's being engaged or exercised vis-à-vis the object of value, whereas the latter cannot).

"Holding" a person responsible thus goes beyond mere evaluative assessment. In the case of the vase smasher, the individual violates a moral norm (amongst others). According to one prominent approach, holding a person morally accountable involves undergoing reactive emotions like resentment and indignation (and perhaps expressing these to the offender)¹². The rough idea on this approach is that moral requirements concern the level of regard and recognition that we ought to accord to each other. Reactive emotions then play the role of tracking the level of regard that other people, in fact, accord to us. Accountability responses are sometimes described as a "backward-looking" on this view given that they are directed towards persons who commit prior wrongdoings.¹³ On "forward-looking" views, by contrast, accountability responses (e.g. sanctions, verbal reproaches, etc.) might be justified on the basis of consequentialist considerations (e.g. their ability to promote socially desirable behavior). Since my topic isn't moral responsibility, I will not explore these different approaches in greater detail here. The point is that, when it comes to "holding" the vase smasher morally responsible for his deed, what is required is that there be *some* form of response that goes beyond basic attributability and negative evaluation. As I

¹¹ I will be primarily focused on instances of *interpersonal* (rather than *intrapersonal*) instances of "holding" responsible. Note, however, that there are instances where we might hold *ourselves* responsible.

¹² Strawson (1962/2003), Wallace (2011; 2019), and Watson (2011).

¹³ Scanlon (2008) develops a view of moral blame which also sees moral requirements as importantly tied to interpersonal relationships. However, on Scanlon's view reactive emotions are not necessary for moral blame.

have indicated, this could potentially involve a number of things: reactive emotions, protest, verbal criticism, sanction, etc. It could also be either “backward-looking” or “forward-looking”.

It is also important to note that, while “blame” is often taken to be the paradigmatic form of moral accountability, there are forms of accountability which do not amount to blame. For instance, there are *legal* forms of accountability, as well as *institutional* or *professional* forms of accountability. The forms of accountability associated with these domains might be adopted vis-à-vis an individual even when the person isn’t also open to moral blame. Consider, for instance, an example from Kauppinen (2018). Say that you decide to drive home after an evening out on the grounds that (a) you don’t feel drunk, and (b) you have been assured by a reliable friend that the drinks you’ve consumed had a very low alcohol content. But now imagine that, in fact, you are over the legal limit, you drive erratically, and are pulled over by the police. In a case like this, according to Kauppinen:

you may well have an excuse from moral sanctions – in the circumstances, it wouldn’t be fair to resent you, say. But you plausibly *don’t* have an excuse from legal sanctions – you should pay your fine. The legal sanction doesn’t imply that you’re morally bad, or that there’s something wrong with your character or your will. (Kauppinen 2018: 9)

If Kauppinen is right, then in a case like this you aren’t an appropriate target of moral blame (e.g. resentment), but you *are* an appropriate target of other forms of accountability (e.g. legal). Kauppinen’s point is that legal accountability responses can be appropriately adopted vis-à-vis an individual even if there’s nothing morally objectionable about the quality of that individual’s character or will (i.e. their “self”).¹⁴ What I will be arguing is that, similarly, an individual like Bill in Case 2 can be the appropriate target of *epistemic* accountability responses even when there’s nothing wrong with the quality of his doxastic self, epistemically speaking.¹⁵

If reactive attitudes like resentment and indignation are common ways to hold an individual *morally* accountable, and if fines and other forms of sanction are common ways to hold an individual *legally* accountable, what might *epistemic* accountability responses look like? As per the above discussion concerning “holding” a person responsible, distinctly epistemic accountability responses, if these exist,

¹⁴ As I interpret Kauppinen’s example, the individual’s decision to drive home fails to disclose an objectionable character or will due to ignorance, not due to impairment. After all, the individual is described as not feeling drunk.

¹⁵ Even though Kauppinen makes the above point regarding the difference between moral and legal accountability, and he also prefers the label “epistemic accountability” over “epistemic blame”, he still seems to understand the former in a way that requires disclosure of an objectionable doxastic self. I will return to this point below (sec. 3).

must somehow go beyond basic attributability and mere negative evaluation. But *do* they exist?

Here I will follow Boulton (2021) and Kauppinen (2018) and take a certain kind of trust modification to be central to epistemic accountability.¹⁶ According to Boulton, what disclosure of an objectionable doxastic self from the epistemic point of view provides is

good reason to suspend [our] presumption of epistemic trust in would-be targets of epistemic blame, at least within some restricted domain, or on some specific matter. (Boulton 2021: 525)

For example, a person might do so by

revising one's intention to epistemically trust the word of another person on a subject matter they seem to be prone to think dogmatically about. (Boulton 2021: 526)

Kauppinen glosses his view in a similar way (albeit without classifying the trust modification as a form of "blame", as Boulton does):

we might say that the basic way of holding someone epistemically accountable is subtracting credibility points from someone. This is perhaps most clearly manifest in the context of *testimony*. When we hold someone to an epistemic norm we endorse and take them to have violated, we no longer take what they say at face value, even if we don't have any ethical doubts about them. We lower our credence in what they say, either in general or about some particular subject matter...and are less willing to partner with them in the project of finding out how things stand regarding it. (Kauppinen 2018:6)

Notice that modifying *epistemic* trust is different from no longer trusting someone because we suspect that they're being dishonest. Thinking that someone is dishonest is compatible with having a very high level of epistemic trust in the person. I might correctly believe that someone is highly knowledgeable about some issue but nevertheless believe that they're unlikely to be honest with me and therefore cease to trust their words. Thus, the kind of trust modification described by Boulton and Kauppinen is *distinctly epistemic*.

This kind of trust modification gives us a promising way of unpacking what it would be to "hold" someone responsible in a distinctly epistemic manner. To illustrate this idea, let's return to Tom in Case 1. Tom's belief meets the condition of basic attributability. It is not something like an itch or a pain, it is not a compulsion, nor was it formed during an altered state of consciousness, etc. It is, we can say, "his" in the relevant sense. Tom's belief is also disclosive of an objectionable doxastic self from the epistemic point of view; it reveals an irresponsible and dogmatic doxastic self insofar as it is maintained in spite of the strong countervailing evidence that is available to him.

¹⁶ I do not think that this is the *only* way that a person might be held epistemically accountable. I only hold that this is one, perhaps central, way in which a person can be held epistemically accountable.

Acknowledging the above doesn't yet get us to "holding" Tom epistemically responsible for his belief, however. For instance, not unlike the detached third-party observer of our vase smasher, we could imagine someone who attributes the belief to Tom and acknowledges that it discloses an objectionable doxastic self but who does nothing more than this. In order to "hold" Tom epistemically responsible for his belief, we must somehow take up further responses towards him; ones that go beyond basic attributability and negative evaluation.¹⁷ Following the above suggestion, the way in which this might be done is by modifying our trust in Tom in certain ways. For instance, we might lower our presumption of epistemic trust in Tom by, say, modifying our intentions to take his words at face value when it comes to the topic of climate change. We might also be less inclined to partner with him in inquiry. Depending on the kind of relationship that we stand in with Tom, we might even explicitly bring this to his attention in certain ways, perhaps by enjoining him to consider the relevant evidence. However, taking up these latter, more direct, forms of accountability will be called for only if we stand in the appropriate relationship with Tom. At the very least, we might just alter our presumption of epistemic trust in the manner described above.¹⁸

These further responses do not seem to follow *automatically* from basic attributability along with negative evaluation. For instance, a person might recognize that Tom's belief is attributable to him and that it violates an epistemic norm, but then abstain from modifying trust in him in the relevant ways. Perhaps this person just didn't think that this particular norm violation was all that important, or maybe they just want to give Tom a second chance, etc. Thus, it's the *actual modification* in trust vis-à-vis Tom which comprises "holding" him epistemically responsible.

But why think this kind of modification actually *does* amount to a genuine way of "holding" a

¹⁷ It is for this reason that purely evaluative notions of epistemic assessment, e.g. William Alston's evaluative concept of epistemic justification (Alston (1989)), do not suffice for accountability responses. According to Alston, evaluation concerns "the possession of certain good-making characteristics." (1989: 97). Just as we can evaluate a knife as "good" or "bad" depending on its having or lacking certain qualities, we could also evaluate a person's beliefs as "good" or "bad" from the epistemic point of view. However, negative epistemic evaluation of a person's beliefs doesn't suffice for holding the person accountable for their beliefs. Thanks to anonymous reviewer for requesting clarification here.

¹⁸ I will argue below that certain trust modifications are also called for in the case of Bill in Case 2. However, there will perhaps be important differences when it comes to the accountability responses that are called for in Case 1 vs. Case 2. For instance, if we have no reason to think that Bill will fail to respond appropriately to new evidence, then we may not be justified in manifesting an unwillingness to partner with him in inquiry. Such an unwillingness may, however, be called for in the case of Tom in Case 1. Nevertheless, as I will argue below, I do think that we can modify our trust in Bill in certain ways, e.g. by ceasing to take his words at face value when it comes to certain topics and issues.

person responsible? Boulton appeals to the fact that this kind of trust modification compromises a certain kind of relationship that we stand in with others, viz. the “general epistemic relationship”. This is the relationship that we stand in with others qua members of epistemic communities, i.e. as testifiers, informants, and partners in inquiry. According to Boulton:

modifying one’s relationship with another person is more significant than simply negatively evaluating them. It is one thing for me to judge that you have done something epistemically remiss, to assign you a bad epistemic grade. It is another for me to modify the intentions and expectations that comprise part of a relationship I stand in with you in response to this shortcoming. Relationships are significant things. (Boulton 2023: 818-819)

Kauppinen also stresses the significance of this kind of interpersonal trust modification, albeit without appealing to Boulton’s relationship-based framework:

Typically, it is bad for you if I don’t take your word for something. It’s even bad for you if I don’t feel let down when you say something that turns out to be false – if I feel that it is to be expected that you get it wrong. We’re not equals in my eyes if I don’t take you seriously as a potential informant. I might be treating you, in effect, like a child. And that’s not a standing that a self-respecting adult wants to have. (Kauppinen 2018: 7-8)

Thus, it appears apt to conceive of the kind of trust modification that we might take up vis-à-vis Tom as a way of holding him responsible in a distinctly epistemic manner.

But in what kinds of circumstances is it appropriate to respond to a person in this way? In the next section, I will turn to this question. I will begin (sec. 3) by looking at a number of recent discussions concerning “epistemic blame” and “epistemic accountability”. As we will see, the existing consensus in the literature is that epistemic accountability responses of the kind described above are only called for when a person (like Tom) discloses an objectionable doxastic self. I will then (sec. 4) challenge this consensus by arguing that we can hold individuals like Bill in Case 2 epistemically accountable.

3. Epistemic Blame, Epistemic Accountability, and Doxastic Self-Disclosure

In the recent literature on epistemic responsibility, a number of authors have been concerned to identify distinctly epistemic forms of interpersonal accountability and/or blame. In developing accounts of these phenomena, most authors incorporate conditions which require disclosure of an objectionable doxastic self. For instance, Boulton’s (2021) relationship-based account of epistemic blame draws inspiration from T.M. Scanlon’s theory of moral blame. Scanlon’s theory explicitly requires disclosure of an objectionable feature of a person’s self:

To blame a person for an action, in my view, is to take that action *to indicate something about the person* that impairs one’s relationship with him or her, and to understand that relationship in a way that reflects this impairment. (Scanlon 2008: 122-123; emphasis added)

As I mentioned earlier, Boulton develops this idea in connection with what he calls the “general epistemic relationship” that we stand in with each other qua testifiers, informants, and partners in inquiry.

According to Boulton, in order for a person to “impair” *this* kind of relationship, they must disclose or reveal some objectionable aspect of their doxastic self:

epistemic blamers tend to judge that someone has been *intellectually irresponsible*, or *intellectually vicious*, or *reckless*, or just plain “*stupid*”. Those are the sorts of things I take the notion of a judgment of general epistemic relationship impairment to unify. Only when an agent modifies their epistemic expectations in a way made fitting by *this* sort of judgment do they count as epistemically blaming others. (Boulton 2023: 818)

Other authors who argue for the existence of a distinctly epistemic form of blame (Brown (2020a), (2020b), Rettler (2018), Schmidt (2021)) either rely on examples in which individuals disclose objectionable doxastic selves, or they explicitly incorporate requirements of the kind defended by Boulton. For instance, according to Brown (2020b),

in the epistemic case, the target of blame is not merely someone whose beliefs violate epistemic norms. For, one could have an excuse for violating an epistemic norm, say because one had strong but misleading evidence that one was not violating it. (Brown 2020b: 399)

An excusing condition is one which renders accountability responses somehow uncalled for or inappropriate, even in spite of a norm violation. Here Brown is suggesting that, when it comes to epistemic blame, individuals like Bill in Case 2 are excused; they are not to be blamed for epistemic norm violations when they have strong (yet misleading) evidence that they aren’t in violation of the relevant norm(s). This, in turn, suggests that epistemic blame tracks certain attributes of the person, e.g. one’s responsiveness to available evidence. Thus, in the context of discussions of “epistemic blame”, incorporation of self-disclosure requirements of the kind mentioned above are common.

Kauppinen (2018) is unique insofar as he reserves the concept of “blame” for the moral domain only.¹⁹ However, Kauppinen agrees that there is a distinctly epistemic form of interpersonal responsibility. He calls the latter “epistemic accountability” rather than “epistemic blame” and, as we saw earlier, he understands it in terms of a certain kind of trust modification. Nevertheless, Kauppinen still appears to incorporate a self-disclosure condition similar to the ones mentioned above. Notice, for instance, the parallels between Kauppinen and Brown when it comes to the topic of epistemic excuses:

an epistemic excuse will be something that shows that in spite of violating an epistemic norm on a particular occasion, the agent is epistemically trustworthy. (Kauppinen 2018: 9)

¹⁹ “What is distinctive of moral norms...is that their violation *prima facie* merits blame in some form – disapproval, resentment, indignation, contempt, or guilt.” (Kauppinen 2018: 5)

The very possibility of an agent's being "epistemically trustworthy" while at the same time violating an epistemic norm suggests that Kauppinen understands epistemic trustworthiness in terms of the attributes of the person rather than the objective accuracy of their doxastic attitudes (or, alternatively, rather than the latter somehow *combined* with the former). Kauppinen continues:

For example, we shouldn't reduce confidence in someone who has been temporarily misled into endorsing a faulty pattern of inference by someone whom they have excellent reason to trust, such as their famous logic professor. We can still make sense of their bad inference as a violation of an epistemic norm rather than merely an evaluative standard: considered in isolation, it would be grounds for reducing epistemic trust. *It would also be grounds for reducing epistemic trust, if it did reflect the subject's belief-forming habits, just in the same way as a morally excused violation would be grounds for blame, if it did reflect the agent's quality of will.* (Kauppinen 2018: 9; emphasis added)

Here Kauppinen is suggesting that epistemic norm violations which are not somehow indicative of bad belief-forming habits are not the appropriate targets of epistemic accountability responses (just as moral norm violations which are not indicative of the quality of a person's will are not the appropriate targets of blame). Kauppinen thus appears to understand epistemic accountability in a way that parallels common views on blame; it only targets epistemic norm violations which reveal or disclose objectionable features of a person's doxastic self.

Kauppinen does admit the possibility of holding epistemically accountable individuals who are in no way blameworthy for *having* objectionable doxastic selves. For example,

someone who grew up in a bizarre, isolated community might acquire very bad belief-forming habits without being to blame for it. Nevertheless, their predictably false beliefs and bad inferences count against them epistemically, and it's appropriate to trust them very little. (Kauppinen 2018: 9)

However, the appropriateness of the trust modifications here still seems to depend upon the bad belief-forming *habits*, according to Kauppinen. This example is thus importantly different from the case of Bill in Case 2. We can stipulate that, in Case 2, Bill's belief-forming habits are flawless, epistemically speaking. The problem in Case 2 is not Bill's habits, but rather the evidence that is available to him, i.e. the "inputs" into his (epistemically superb) personal-level processes of belief maintenance and regulation. Given the textual evidence cited above, I am inclined to read Kauppinen as denying that individuals like Bill are the appropriate targets of epistemic accountability responses.²⁰ In the next section, I will argue

²⁰ In addition to the evidence cited above, Kauppinen also says of his illustrative example of an individual who is appropriately held epistemically accountable that "there is something wrong with *the way he processes information and evidence in forming beliefs and making claims*, at least on this occasion."

that this verdict is mistaken: Bill can be appropriately held epistemically accountable for his beliefs concerning human-induced climate change.

4. Holding Epistemically Responsible without Disclosure of an Objectionable Doxastic Self

I would now like to argue that Case 2 represents an instance in which all of (1)-(3) are met. Let us consider each of these conditions in turn. First, Bill's belief seems to meet the condition of basic attributability. Not unlike Tom's belief, Bill's belief is not akin to an itch or a pain, it is not a compulsion, nor was it formed while Bill was in an altered state of consciousness, etc. Bill is, of course, the victim of circumstances that are beyond his control. We shouldn't blame him for being in the predicament that he finds himself in. Nor, I think, should we blame him for the false beliefs that he forms as a result of being in this predicament (although recall that Kauppinen would argue that we shouldn't "blame" Tom in Case 1 either; we should hold him epistemically accountable). Nevertheless, it doesn't seem that being in this kind of environment jeopardizes basic attributability; Bill's belief is still "his" in the sense that opens up the possibility of accountability responses. Thus, (1) is met in Case 2.

Bill's belief also does not disclose or reveal anything objectionable about his doxastic self when the latter is assessed from the distinctly epistemic point of view. Bill is interested in forming true beliefs about the relevant subject matter, and he frequently inquires into the topic by doing research and seeking out evidence. In other words, Bill isn't dogmatic, biased, lazy, or hasty in believing what he does; his belief does not reveal an objectionable feature of his doxastic self. Thus, (2) is met in Case 2.

Finally, (3) is met in Case 2 as well. Once we acknowledge that "holding" a person epistemically responsible can consist in a kind of interpersonal trust modification (which need not be classified as "blame"), then it seems that we can "hold" Bill epistemically responsible for believing what he does. For instance, it would be entirely appropriate to not take Bill's words at face value when it comes to the topic of climate change. We might also decide to not seek him out as a testimonial source of information when it comes to this subject matter. Depending on the exact nature of our relationship with Bill, we might even hold him accountable in more direct ways, e.g. by enjoining him to consider certain evidence which he hasn't yet seen. However, even if we don't do this, we can still hold him accountable in the ways described above.

(Kauppinen 2018: 6; emphasis added). Immediately following this, Kauppinen asks: "Is there some kind of negative consequence that we could impose *that would target precisely this aspect of his functioning...*?" (2018:6; emphasis added).

The nature of Case 2 may lead one to wonder how we might be able to take up these responses to Bill if *we're* not also members of his community. Here I'd like to note two points. First, there's an expansive notion of "testimony" which incorporates more than just direct, face to face communication between two individuals. For instance, Bill might communicate his views through many different avenues: emails, letters, text messages, videos, blog posts, social media posts, etc. Thus, we can imagine that Bill's community isn't so shut off from the outside world that he doesn't have access to certain ways of communicating his thoughts which are capable of reaching an outside audience. As a result, individuals outside of Bill's insulated community can cease to take Bill's claims at face value when it comes to the topic of climate change. Second, we can also imagine a scenario in which Bill *leaves* his insulated community and enters the outside world. At some time after Bill leaves his community but before he confronts and properly digests evidence which shows that his existing views on climate change are false (something which may take some time), individuals can still hold Bill epistemically responsible in the manner discussed above.

I thus think that Case 2 gives us a clear example of an instance where all of (1)-(3) are met. While the case is fictional, I think it parallels many realistic examples in the actual world. Imagine, for instance, individuals who grow up in families, communities, or societies where dogmatic allegiance to certain views (religious, political, ethical, etc.) is typical. In such cases, there may be very real limits on the kinds of evidence that individuals have access to. Moreover, even if there are certain avenues via which individuals could uncover more evidence, they may (through no fault of their own) lack the wherewithal to pursue these avenues. There may also be cases where individuals blamelessly find themselves in "epistemic bubbles" and, as a result, maintain false beliefs about various topics.²¹ I grant that, in any given case, whether or not (2) is met will be a difficult interpretive matter. Nevertheless, I do think it will be met in many actual cases where (1) and (3) are also met.

In the remainder of this section, I would like to consider and respond to three different objections to the argument that I have just presented. As we will see, all three of these objections involve denying that (3) is met in Case 2, albeit for different reasons. Here are the three objections:

1. (3) is not met in Case 2 because modifying epistemic trust never, *by itself*, constitutes a way of holding a person responsible, even in a case like Case 1. More is required in order to "hold" a person responsible.

²¹ For discussion see: Nguyen (2020).

2. (3) is not met in Case 2 because it *isn't* appropriate to modify our epistemic trust in Bill in Case 2. Modifying epistemic trust in a person *can*, by itself, constitute a way of holding a person responsible, and this is not an appropriate way to respond to Bill in Case 2.
3. (3) is not met in Case 2 because, even though epistemic trust modifications are called for in *both* Case 1 and Case 2, these modifications only constitute a way of holding a person responsible in the former case, not the latter.

I will answer each of these objections in turn.

Objection 1

A defender of the first objection would disagree with authors like Boulton and Kauppinen and argue that modifying epistemic trust in a person doesn't, by itself, constitute a way of "holding" them responsible. Even in a case like Case 1, more is required, e.g. reactive emotions or overt criticism/censure. Here I'd like to mention two things. First, recall that I am not committing myself to the view that this form of response should be classified as "blame", even in cases like Case 1. I have only committed myself to the view that this kind of trust modification constitutes a way of holding a person accountable. Holding someone accountable can be an active, engaged form of response vis-à-vis an individual even when it doesn't amount to blame. Second, I think both Boulton and Kauppinen's claims about the significance of this kind of interpersonal trust modification are apt. Whether or not we accept the details of Boulton's relationship-based account, it does seem correct to say that our standing vis-à-vis one another qua testifiers and informants is a significant social relationship. Insofar as the kind of interpersonal trust modification discussed above negatively affects our standing vis-à-vis others qua testifiers and informants, it can be seen as a socially significant response; one which provides the means whereby epistemic norms are interpersonally enforced.

Objection 2

A defender of the second objection would argue that it actually *isn't* appropriate to modify our level of epistemic trust in Bill in Case 2. Perhaps this is an appropriate response in Case 1, but not in Case 2. While this response has some merit, I do not think it succeeds. Specifically, I think a lot hinges here on the specific interpersonal dynamics in different cases. Imagine, once again, that Bill manages to leave his insulated community and enter the wider world. If we became Bill's *friend* or *teacher*, or if Bill *chose* to engage with us in conversation regarding the relevant topic, then it seems that what we should do is not suspend our presumption of epistemic trust in Bill. Rather, we should try to bring to Bill's attention any relevant evidence which he has yet to consider. I grant that doing so would not involve lowering epistemic trust in Bill. Nevertheless, I do not think this threatens my argument. First, notice that taking up

these alternative, more direct, forms of response vis-à-vis Bill (e.g. calling his attention to further evidence) might *themselves* be ways of holding Bill accountable to an epistemic norm. Once again, these forms of interaction need not be construed as blame responses. Second, for people who do *not* stand in these specific relationships with Bill (friend, teacher, etc.) it does not seem that it would be appropriate for them to confront Bill about his beliefs, or to present him with new evidence. To do so would be presumptuous and overbearing. However, in lieu of these more direct responses, what are such individuals left to do when it comes to Bill? I think the only plausible option here is for them to lower their epistemic trust in Bill when it comes to certain topics and issues, albeit without blaming him.

Objection 3

A defender of the third objection would argue that trust modifications are called for in *both* Case 1 and Case 2. However, when we modify trust in Bill in Case 2, this doesn't amount to a form of interpersonal accountability. This is different from the first objection. A defender of this last objection is willing to grant that, when we modify trust in Case 1, it does indeed constitute a genuine form of interpersonal accountability. However, when we modify trust in Case 2, it doesn't. Thus, according to this view, only when our epistemic trust modifications are also underwritten by a recognition that an objectionable doxastic self has been disclosed does the trust modification actually constitute a way of holding a person responsible. As we saw earlier, this is the prevailing view when it comes to the recent literature on "epistemic blame" and "epistemic accountability".

Defenders of this third objection owe us an account of what makes epistemic trust modifications rise to the level of accountability response, one which allows them to sort cases in the way that they want. Usually, that story is filled in by invoking a condition which requires disclosure of an objectionable self. However, appealing to this alone would be question-begging in the present context; I have given an argument for thinking that the trust modifications in Case 2 constitute accountability responses in spite of the fact that Bill does not disclose an objectionable doxastic self. What else might a defender of this third objection say?

Here's one possibility: There are various ways in which our relationships with others can be modified or altered without these alterations also rising to accountability responses. Perhaps the trust modifications that we can adopt towards Bill in Case 2 are similar to this. For instance, imagine two friends who simply drift apart. Say that there was no incident or argument which precipitated the modification of their friendship; maybe one of the friends just moved away and they lost touch. Also,

consider cases where we might alter our relationship with a person after they have become unpredictable due to some form of mental illness or impairment. Examples like this illustrate the possibility of modifying our relationship with another person where this does not amount to holding them responsible for some norm violation. The suggestion, then, is that the trust modifications that can be adopted vis-à-vis Bill in Case 2 are somehow analogous to cases like these.²²

I think we can immediately rule out the possibility that our trust modifications vis-à-vis Bill resemble cases where we modify our relationship with another due to something like mental illness or impairment. Recall that, as I've already argued, Bill's belief is "attributable" to him; it is not something which fails to express or disclose "who he is" qua believer due to some internal issue which threatens his rational agency. On the contrary, as I have stipulated, Bill is fully rational and he is manifesting a flawless "doxastic self" from the epistemic point of view. I agree that there are ways of altering our level of trust in the words of another which might be characterized as "mere differential reliance". These could be seen as interpersonal analogues of the kinds of trust modifications that we can adopt towards objects and artefacts, e.g. malfunctioning clocks or thermometers. When we modify our trust in the words of another in this way, we have not held the person accountable; perhaps we have just realized that they are talking in their sleep, or are under hypnosis, or are fully in the grip of a paranoid delusion, etc.²³ However, this kind of stance is entirely uncalled for when it comes to Bill in Case 2. As I have indicated, this is the kind of stance that we might adopt towards someone when their rational capacities are somehow inactive or compromised. However, as I have noted, Bill is not impaired, and his rational capacities are fully intact.²⁴

This leaves us with examples like the two friends who have drifted apart. Once again, this is not the right way to think about Bill in Case 2. Notice two important features of Case 2 which distinguish it from cases like this. First, Bill *still stands in* the relevant relationship with others, viz. the testimonial

²² I would like to thank an anonymous referee for raising this challenge.

²³ Cf. Craig (1990)'s helpful distinction between "informants", on the one hand, and mere states of affairs with evidential import, on the other. According to Craig, "the distinction is that between a person's *telling* me something and my being able to tell something from observation of him" (Craig 1990: 35; emphasis added). As I understand the distinction, "informants" are individuals who are capable of participating in cooperative speech acts, e.g. acts of "telling". Bill is obviously still capable of this. Individuals who are, say, under hypnosis are not (even though we might still treat their utterances as having evidential import).

²⁴ Another way to put the central point of this paragraph is that it is inappropriate to take up vis-à-vis Bill what Peter Strawson called an "objective" attitude. Strawson thought this kind of attitude can be taken up towards an individual when they are "warped or deranged, neurotic, or just a child" (Strawson 1962/2003: 66). It is inappropriate to view Bill in this way, even for individuals who are not members of his community.

relationship. This should be especially clear once we acknowledge the “expansive” conception of testimony mentioned earlier. Second, as Scanlon (2008) has noted, in the case of two friends who have drifted apart, there is no “violation of the standards of friendship” (2008: 135). Because of this, the alteration of the relationship in cases like this doesn’t involve one person holding another responsible. According to Scanlon, the latter occurs when

*one party, while standing in the relevant relation to another person, holds attitudes...that are ruled out by the standards of that relationship, thus making it appropriate for the other party to have attitudes other than those that the relationship normally involves. (Scanlon 2008: 135; emphasis added)*²⁵

This is not what we have in the case of two friends who have drifted apart. First, they no longer stand in the relevant relationship. Second, the dissolution of the friendship wasn’t precipitated by one of the friend’s holding attitudes that are ruled out by the standards of friendship. As Scanlon notes, “[f]riendship would be an oppressive relationship if it had to last forever, no matter what” (2008: 135). However, my suggestion is that this *is* what we have in the case of Bill in Case 2. First, as a rational and responsible epistemic agent, Bill still stands in the relevant relationship with others. What would it be for a person to no longer stand in this kind of relationship with others? One type of case might involve individuals who, because of various forms of incapacity or impairment, are no longer capable of participating in testimonial interactions or engaging in certain speech acts. Another type of case might involve individuals who, say, have decided to live their lives in solitude, abstaining from testimonial interactions with other people. Bill clearly does not fit into either of these camps; he still stands in the relevant relationship with others. Second, by maintaining false beliefs about human-induced climate change, Bill holds attitudes that make it appropriate for others to respond to him in certain ways, e.g. by lowering their epistemic trust in him. Modifying trust vis-a-vis Bill in this latter way involves taking on attitudes “other than those that the relationship normally involves” (Scanlon 2008: 135), where the relevant relationship is the testimonial relationship. This implies that Bill holds doxastic attitudes that violate the standards governing testimonial interactions, where “testimony” is construed expansively so as to include more than just direct, face to face communication between two people.

But what about Bill’s praiseworthy doxastic self? Doesn’t Bill’s doxastic self insulate him from negative accountability responses vis-à-vis his false beliefs concerning human-induced climate change?

²⁵ While Scanlon’s focus is on moral blame, his account will be helpful when it comes to distinguishing Case 2 from the example of two friends who have drifted apart.

No, it does not. Recall Kauppinen's example of the blameless drunk driver. That person's laudable "self" excused them from *blame*, not from negative accountability responses. As Kauppinen correctly noted, that person is still the appropriate target of legal accountability responses. My suggestion is that, similarly, Bill can be held epistemically accountable for his doxastic attitudes. Just as erratic driving can call for certain forms of accountability even when the driver shows no ill-will or reckless disregard, believing something false can call for certain forms of accountability even when the believer shows no objectionable doxastic self from the epistemic point of view. Bill's praiseworthy doxastic self thus does not excuse him from being held epistemically accountable; it perhaps excuses him from epistemic *blame*, but, as I have stressed, we should be careful not to conflate blame with accountability.

I have argued that there are cases where all of (1)-(3) are met. I do not mean to suggest that disclosure of an objectionable doxastic self from the epistemic point of view is entirely irrelevant when it comes to holding individuals epistemically responsible. For instance, there might be certain responses that are called for in cases like Case 1 which are not called for in cases like Case 2.²⁶ I'm open to the possibility that there are a range of accountability responses in the epistemic domain, perhaps even including a distinctly epistemic form of blame. My main objective was to establish that there are instances where person is legitimately held epistemically responsible even when they have revealed nothing objectionable about their doxastic self.

4. Conclusion

I would like to conclude by connecting the forgoing discussion with some more general observations about the study of epistemic normativity. Recently, a number of authors have stressed that our social-accountability practices might provide important insights about the normative status of the epistemic domain. For instance, according to Schmidt (2021),

in order to get clear about the normative significance of evidence, we need to think about the reactive attitudes that are appropriate towards violations of purely epistemic norms...The normativity of evidence can become intelligible by understanding this practice. For the reactive attitudes within our epistemic practice reveal the normative significance of purely evidential considerations. (Schmidt 2021: 3)

Kauppinen (Forthcoming) also puts forward a similar claim. Additionally, Goldberg (2018) defends the claim that epistemic norms themselves are grounded in/explained by our social practices of interpersonal accountability rather than the other way around. The unifying idea being pursued by such authors is that

²⁶ See note 18.

our social practices of interpersonal accountability might be capable of shedding light on the nature of epistemic normativity itself. These practices might, for instance, bear on the question of whether or not the epistemic domain comprises an independent normative domain, and whether or not the content of epistemic norms incorporates objectivist/externalist notions (e.g. knowledge or truth).

If these authors are correct in thinking that our social-accountability practices are capable of shedding light on the nature of epistemic normativity itself, and if I'm correct in thinking that there are cases where all of (1)-(3) are met, then theorists of epistemic normativity should be concerned with more than just accountability responses which track objectionable features of a person's doxastic self.

Acknowledgments

I would like to thank two anonymous reviewers and the editors at the journal *Ergo*.

References

- Alston, William. (1989). *Epistemic Justification: Essays in the Theory of Knowledge*. Cornell University Press.
- Boult, Cameron. (2021). "There is a Distinctively Epistemic Kind of Blame". *Philosophy and Phenomenological Research*, 103(3):518–534
- (2023). "The Significance of Epistemic Blame". *Erkenntnis*, 88: 807-828.
- Brown, Jessica. (2020a) "Epistemically Blameworthy Belief". *Philosophical Studies*, 177:3595-3614.
- (2020b) "What is Epistemic Blame?" *Nous*, 54:2, 389-407.
- Coates, D. Justin. (2020). "The Ethics of Blame: A Primer" in *The Ethics of Belief and Beyond: Understanding Mental Normativity*, ed. Schmidt, S. & Ernst, G. New York: Routledge.
- Craig, Edward. (1990). *Knowledge and the State of Nature*. Oxford University Press.
- Dewey, John. (1891). *Outlines of a critical theory of ethics*. Ann Arbor, MI. Register Publishing Company.
- Frankfurt, Harry. (1971). "Freedom of the Will and the Concept of a Person". *The Journal of Philosophy*, 68(1), 5–20.
- Flores, Carolina and Woodard, Elise. (Forthcoming). "Epistemic Norms on Evidence-Gathering" *Philosophical Studies*.
- Friedman, Jane. (Forthcoming). "Zetetic Epistemology" in *Towards an Expansive Epistemology: Norms, Action, and the Social Sphere*, eds. B. Reed & A.K. Flowerree. Routledge.
- Goldberg, Sanford. (2018). *To the Best of Our Knowledge*. Oxford University Press.

- Hieronymi, Pamela. (2008). "Responsibility for Believing". *Synthese*, 161, 357-373.
- Kauppinen, Antti. (2018). "Epistemic Norms and Epistemic Accountability" *Philosophers Imprint*. 18(8).
 ----- (Forthcoming). "The Epistemic vs. the Practical" *Oxford Studies in Metaethics* vol. 18.
- McHugh, Conor. (2013). "Epistemic Responsibility and Doxastic Agency". *Philosophical Issues*, 23, 132-157.
- Nguyen, C. Thi. (2020). "Echo chambers and epistemic bubbles". *Episteme*, 17(2), 141-161.
- Rettler, Lindsay. (2018). "In Defense of Doxastic Blame". *Synthese*, 195, 2205-2226.
- Scanlon, Thomas. (1998). *What we Owe to Each Other*. Harvard University Press.
 ----- (2008). *Moral dimensions: Permissibility, meaning, blame*. Harvard University Press.
- Schmidt, Sebastian. (2021). "Epistemic Blame and the Normativity of Evidence". *Erkenntnis*, <https://doi-org.proxymib.uits.uu.edu/10.1007/s10670-021-00430-9>
- Shoemaker, David. (2015). "Ecumenical Attributability". In R. Clarke, M. McKenna, & A. Smith (Eds.), *The nature of moral responsibility: New essays*. New York: Oxford University Press.
- Smith, Angela. (2005). "Responsibility for Attitudes: Activity and Passivity in Mental Life". *Ethics*, 115, 236–271.
 ----- (2015). "Responsibility as Answerability" *Inquiry*, 58(2), 99-126.
- Sripada, Chandra. (2016). "Self-expression: A Deep Self Theory of Moral Responsibility". *Philosophical Studies*, 173(5), 1203–1232.
- Strawson, Peter. (1962/2003). "Freedom and Resentment". *Proceedings of the British Academy* 48, 1–25.
 Reprinted in Gary Watson (ed.) *Free Will*. Oxford: Oxford University Press, 72–93.
- Reis-Dennis, Samuel. (2018). "Responsibility and the Shallow Self". *Philosophical Studies*, 175, 483-501
- Taylor, Charles. (1976). "Responsibility for Self". *The identities of persons*, 281-299. University of California Press.
- Wallace, R. Jay. (2011). "Dispassionate Opprobrium: On Blame and the Reactive Sentiments." In R.J. Wallace, R. Kumar, & S. Freeman (Eds.), *Reasons and Recognition: Essays on the Philosophy of T.M. Scanlon* (pp. 348-372). Oxford University Press.
 ----- (2019). *The Moral Nexus*. Princeton University Press.
- Watson, Gary. (1975). "Free Agency". *The Journal of Philosophy*, 72, 205–220.
 ----- (2004). "Two Faces of Responsibility" in *Agency and Answerability*. Oxford University Press.
 ----- (2011). "The Trouble with Psychopaths". In R.J. Wallace, R. Kumar, & S. Freeman (Eds.), *Reasons*

and Recognition: Essays on the Philosophy of T.M. Scanlon (pp. 307-331). Oxford University Press.