



THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Believing as we ought and the democratic route to knowledge

Citation for published version:

Chrisman, M 2020, Believing as we ought and the democratic route to knowledge. in S Schmidt & G Ernst (eds), *The Ethics of Belief and Beyond: Understanding Mental Normativity*. Routledge Studies in Epistemology, Routledge, London. <https://doi.org/10.4324/9780429284977>

Digital Object Identifier (DOI):

[10.4324/9780429284977](https://doi.org/10.4324/9780429284977)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Peer reviewed version

Published In:

The Ethics of Belief and Beyond

Publisher Rights Statement:

This is an Accepted Manuscript of a book chapter published by Routledge in The Ethics of Belief and Beyond: Understanding Mental Normativity on 2 April 2020, available online: <https://www.routledge.com/The-Ethics-of-Belief-and-Beyond-Understanding-Mental-Normativity/Schmidt-Ernst/p/book/9780367245504>

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.



Believing as We Ought and the Democratic Route to Knowledge¹

Matthew Chrisman

Philosophy
University of Edinburgh
3 Charles Street
Edinburgh EH111RP
United Kingdom

ORCID iD <https://orcid.org/0000-0002-2886-5829>

Abstract: In the attempt to understand the norms governing believers, epistemologists have tended to focus on individual belief as the primary object of epistemic evaluation. However, norm governance is often assumed to concern, at base, things we can do as a free exercise or manifestation of our agency. Yet believing is not plausibly conceived as something we freely do but rather as a state we are in, usually as the mostly automatic or involuntary result of cognitively processes shaped by nature, bias, and ideology. In this paper, I sketch a response to this tension. This response is based on rejecting the traditional theoretical focus on an individual's particular beliefs as the primary object of normative epistemic evaluation. If we shift our focus from the particular beliefs of individuals to the community and its information managing practices, we may lessen the tension between norm governance and automaticity and involuntariness in a way that construes autonomous cognitive agency as a resultant of rather than a precondition for our norm-governed epistemic sociality.

1. Introduction

Here is one way to think of the theory of knowledge and some of its recent history: Knowledge is a relation an individual can stand in to true propositions, and epistemology's job is to figure out what this relation is like. To do so, we start out with the platitude that knowing *p* requires believing *p* truly, and then we go off searching for the extra ingredient that makes the difference between true belief and knowledge.

In the heyday of naturalized epistemology, it was tempting to think this extra ingredient could be some nonnormative (usually modal) feature of the belief, such as its reliability, sensitivity, or safety. But these modal properties come in degrees, so we might need a normative relation after all (something like being as reliable, sensitive, or safe as the relevant kind of belief *ought* to be). Moreover, another strand in the theory of knowledge has long stressed the connection between knowledge and normative

¹ For helpful feedback on this paper, I'd like to thank Graham Hubbs, Sophie Keeling, Sebastian Köhler, Sebastian Schmidt, Nick Treanor, and Sam Wilkinson, as well as the audience at the Ethics of Belief workshop in Erlangen in 2017.

standards for belief such as being reasonable, rationally supported by the evidence, warranted, and/or the manifestation of intellectual virtue.

This encouraged a normative approach to the theory of knowledge, which is my starting point here. Assuming knowledge requires true belief, epistemology's job would then seem to be one of figuring out what *normative* standard an individual believer must meet in order for her particular true belief to count as knowledge. And coming to this via the "extra ingredient" idea has encouraged a lot of contemporary epistemology to follow a methodology of contrasting cases. That is, we consider a case of some individual believer who we intuitively think knows that *p* and contrast this with cases of some other individual believer in the same true proposition who we intuitively think doesn't know that *p*; then we try to determine what the normatively relevant difference is between the cases.

On this backdrop, my aim in this paper is twofold. Ultimately, in section 6-7, I want to add contribute to the motivation to and development of an alternative approach to the theory of knowledge, which as emerged in recent years, centering on the social role of knowledge in communities of interdependent believers with diverse perspectives on the world, rather than a methodology of contrasting cases about individual believers. It's not that I think intuitions about those cases are worthless, but I suspect they are symptoms of a deeper albeit more implicit understanding of knowledge developed through our lived experience as members of epistemic communities, where various normative standards are at play in regulating our economy of shared knowledge. If we can make this social epistemological understanding more explicit, I think we will have a better shot at an illuminating theory of knowledge.

Initially in section 2-5, in order to explain why I think the more individualistic approach to the theory of knowledge characteristic of traditional epistemology is problematic, I want to consider an important objection to the normative approach to the theory of knowledge, pursued through the methodology of contrasting cases. Although this will need to be refined below, the objection is basically that, if we're going to reasonably say (in the way characteristic of genuine normative evaluation) that you ought to do something, then it must be up to you whether you do it; but what you believe isn't (typically or perhaps ever) up to you, the individual believer. So there is a false presupposition of *can* in the idea that there are ways individuals *ought* to believe that would make the difference between an individual knowing and not knowing.

If you're an epistemologist engaged in developing a theory of knowledge, especially if you are pursuing the normative approach, then you are probably already thinking of ways to respond to this objection. Maybe there's a kind of autonomous activity closely related to belief that underwrites our holding individual believers to a standard of what they epistemically ought to believe despite much belief being automatic, involuntary, or otherwise outside the scope of autonomous action. Or maybe there's a distinctive "purely evaluative" kind of normativity that doesn't presuppose any kind of autonomy, and epistemic norms should be assimilated to that kind of normativity. In my experience, however, making out such responses is not as easy as it looks. So, in the central parts of this paper, I want to consider some individualistic ways of responding to this 'ought' implies 'can' objection that I find attractive in various respects but then

explain why I don't find any of them fully satisfactory. The take-home message will be that they all leave out a kind of autonomous agency we enjoy only through our sociality.

In political philosophy, this idea of a kind of autonomous agency enjoyed only through sociality is often attributed to Rousseau. I won't get into any of the details of Rousseau's political morality here, but my organizing intuition is that recognizing a kind of *intellectual* autonomy based in our sociality (alongside Rousseau's notion of a kind of *political* autonomy based in our sociality) can help us to make better sense of normative standards a belief must meet in order to count as knowledge, even while granting that most individual belief is automatic, involuntary, or otherwise outside the scope of autonomous action.

Even with recently increased interest in social epistemological issues, the standard (conservative) order of approach is to start with individual knowers and try to figure out what makes them knowers; then on this basis we seek to answer social epistemological questions about things like the nature of knowledge through testimony, highly complex group knowledge, knowledge extended through technology, and social/political structures needed to recognize, protect, and promote knowledge. I am interested in all of these questions, but I suspect the order of approach needs to be reversed: first social epistemology, then individual epistemology. And I intend my suggestion that there is a kind of intellectual autonomy based in our sociality to provide important groundwork for reversing it, thereby suggesting an alternative methodology for the theory of knowledge.

2. The 'Ought' Implies 'Can' Objection to Norms of Belief

If we are pursuing the theory of knowledge individualistically and normatively, then we assume that there are true propositions one ought to believe and true propositions one ought not to believe; and this 'ought' is understood to be distinctively *epistemological* – i.e. not moral, prudential, or anything else. Believing as one ought, in this sense, is a normative standard one must meet in order to count as knowing.² This is why one central way to criticize someone epistemologically is to suggest that they believe something that they (epistemologically) ought not to believe.

The main challenge to this idea, raised prominently by Williams (1973) and Alston (1988), relies on the assumption that genuine normative demands on someone to ϕ presuppose that ϕ -ing is the sort of thing they can decide to do, or voluntarily do, or freely exercise their will in doing, or something like that. I'll consider challenges to this assumption below, and we should note that the problem looks slightly different

² This is consistent with the idea that, for each true proposition, there is no all-things-considered requirement to know it. I think there are all-things-considered requirements to know some things; however, the kind of normativity I'm focussing on here is domain-specific and sometimes appealed to in making sense of the difference between knowledge and other true belief. The idea is to consider things people believe and ask whether they ought to believe those things or something different, but only from the perspective of determining whether those beliefs are instances of knowledge.

depending on how one refines it; but the initial point is that believing doesn't seem to be like autonomously doing something.

For one thing, believing is not typically conceived as something one *does* but rather a way that one *is*. That's because belief is typically and for good reason assumed to be a mental state rather than a mental action or activity.³ But there's a deeper worry here based in a picture of our mental lives that is radically at odds with the (caricature) picture sometimes thought to be implicit in a lot of traditional epistemology. The supposedly traditional picture is one whereby the paradigmatic cognitive activity is deliberation, which is conceived as a process of considering one's evidence, or what one takes to be reasons, for and against a proposition *p* in the course of deciding whether to believe *p*. On this picture, although belief is stative, it is understood to be the product of something active: deliberation, i.e. consciously weighing up and responding to reasons. But as Williams and Alston suggest and as much recent work in cognitive psychology and behavioral economics seems to confirm, a more realistic picture of our mental lives is one according to which most beliefs are clearly *not* the product of anything like conscious weighing up and responding to reasons.⁴

It may be tempting to respond by arguing that the mental states reached in ways other than deliberation are proto-beliefs or posited degrees of confidence; or maybe they are instances of "animal belief" posited as representational states useful for explanations some behavior but distinguishable from the fully human beliefs that are the subject of traditional epistemology and the possible objects of normative evaluation. For example, De Sousa argues that we should separate the kind of subjective confidence that figures in prediction and explanation of behavior, both animal and human, and the mental states classically called "belief". He writes, "The Bayesian theory provides us with a way of describing the mechanism of nonverbal deliberation in humans and other animals. At a more sophisticated level, our own reliance on language adds incalculably to the scope and complexity of our deliberations, by providing us with a stock of sentences accepted as true. This is what the classical notion of belief is designed to capture." (1971: 57-58) And Evans influentially distinguished between "information states" and "beliefs," writing, "It is as well to reserve 'belief' for the notion of a far more sophisticated cognitive state: one that is connected with (and, in my opinion, defined in terms of) the notion of judgement, and so, also, connected with the notion of reasons." (1982: 124)

However, these "levels of mentality"⁵ seem to me to cut across possible objects of knowledge or at least objects of what we should be trying to understand when developing a theory of knowledge. We could of course reserve the word 'belief' for mental states only sophisticated and linguistically capable reasoners acquire through

³ I make this case in more detail in Chrisman (2012, 2017). See also Steward (1997) and Hunter (2001). for important defences of the idea that beliefs are states a person is in rather than states inside of a person.

⁴ I have in mind the vast literature stemming from Kahneman and Tversky (1973). But see also Mercier and Sperber (2017) for a picture of reasoning that doesn't depend on the system 1 vs. system 2 distinction and still portrays human belief formation as mostly unreasoned.

⁵ Stevenson (2002) argues that there are six different levels of belief and that it is mainly an uninteresting verbal question whether we reserve the term 'belief' for the highest levels. I'm inclined to agree and think that belief at each level is a possible object of knowledge. So the theory of knowledge approached as a normative project will need to explain how normative standards can be applied to all of these different kinds of mental states, whether or not we call them all "belief".

deliberation on reasons, but then we will just need another word for the mental state which, when true, can count as knowledge provided that its possessor meets some normative standard. So I'm inclined instead to recognize lots of different levels or kinds of belief. Moreover, even amongst the kinds of belief only accessible to sophisticated and linguistically capable humans, it remains highly implausible that most of them are the product of deliberation about reasons.

To see why, consider the following. Most beliefs about our perceptible environment, even sophisticated ones supposedly available only to linguistically capable beings, appear to be automatic and unchosen. (E.g., ask yourself whether you deliberated over whether to believe there are words on this page.) Indeed, as Hume (1739-40: I.iv) recognized long ago, it is actually quite important for our survival for most ordinary beliefs, especially those about useful things in our immediate environment, to be shielded from the slow and highly fallible processes of human deliberation.⁶ A further part of this picture is that beliefs, especially those about humdrum though unperceived facts, are stereotypically formed via intuition or testimony without anything like explicit deliberation. (E.g. ask yourself whether you at some point chose to believe that the currency of Mexico is the peso and did so for reasons you weighed up or even considered). Forming beliefs automatically and intuitively rather than based on deliberation is a cognitive necessity; without it we'd commonly suffer paralyzing processing overloads. And, to pick out a different kind of example, contemporary political epistemology and cognitive psychology is increasingly coming to understand how many of our beliefs are strongly influenced by typically unconscious heuristics and biases. Moreover, what we regard as evidence for and against controversial propositions seems to be tacitly though strongly framed by background ideologies on which we rarely reflect explicitly.

Nothing in this picture of human cognitive psychology implies that we *never* form beliefs by deliberating on evidence or reasons, nor does it entail that it is impossible to decide to believe, or to believe voluntarily, or to exercise one's will in believing. Of course, we philosophers are very familiar with being unsure about some argument or objection, thinking about it for a while, discussing it with colleagues and students, and finally deciding what should be believed. And there may be special cases of belief that are susceptible to explicit decision. For example, Velleman (1985) argues that deciding to perform some action, such as to take a walk, is identical to forming the belief that one will perform the action; and surely in normal circumstances one can voluntarily decide to take a walk, and indeed exercise one's will in so deciding. So, if Velleman is right (which I don't mean to assert that he is), then at least that kind of belief is subject to autonomous decision. In a different vein, McCormick (2011, 2015) has argued that, in cases where the evidence isn't conclusive, one can act on practical reasons in freely choosing to believe, e.g. that there is a God or that one's spouse is innocent. So, if McCormick is right (which I don't mean to assert that she is), it would appear that some beliefs can be voluntarily chosen.

⁶ Moreover, the psychological law-governed and predictable effects of having a belief seem to be more correlated with beliefs one does not consciously deliberate to or about. Compare Mandelbaum (2016) for a discussion of implicit biases and an interpretation of psychological data supporting this suggestion.

To a proponent of the previous picture of human cognitive psychology, however, these kinds of cases will at best look like exceptions that prove the rule, not paradigm cases of belief formation. And, if that's right, then epistemologists pursuing the normative approach to the theory of knowledge still face a challenge stemming from the apparent involuntariness of belief. We want the theory of knowledge to apply to all knowledge, or at least all paradigm cases. Yet, in a lot of paradigm cases, it doesn't look as if we have the kind of voluntary control, freedom, autonomy, or whatever that is presupposed when we hold individual believers to normative standards about what they ought (epistemologically) to believe.

As I see things, this 'ought' implies 'can' objection is the fundamental challenge to a normative approach to the theory of knowledge, which is one of the dominant approaches in contemporary epistemology.⁷ Next, I want to consider three responses to the challenge, all of which have some attraction but none of which is fully satisfactory, in my view. The first response can be classified as seeking to identify a kind of genuine epistemological norm or 'ought' which doesn't however imply individual believers have autonomy over what they believe. The second and third responses can be classified as seeking to identify a kind of autonomous activity individuals engage in that is intimately related to belief but different from paradigm autonomous action.

3. Norms of Belief as Evaluative rather than Prescriptive

The first response to the objection that I want to consider rejects the implication from 'ought' to 'can'. Of course, we know that the word 'ought' is used to make probabilistic claims that have nothing to do with things like decision, voluntary control, or autonomy. The weather reporter's suggestion that the storm ought to hit land by nightfall doesn't impute any agency to the storm, and it is not a normative evaluation of the storm in the way that claims about what people epistemologically ought to believe are normative evaluations of believers. So a proponent of this response needs a genuinely normative sense of 'ought' that nonetheless doesn't imply 'can' in the sense of being something up to the person evaluated. Is there such a sense of 'ought'?

We might find something like this sense of 'ought' in the suggestion that, while paradigmatic norms of action are prescriptive, some other norms are "evaluative" rather than prescriptive.⁸ What this means is that the function of these norms in our thinking is to set a standard for evaluating something rather than prescribing action. However, it can't be any old standard that is an evaluative norm. Being as tall as George Clooney is a standard we can use to evaluate heights, but this isn't (all by itself) a normative standard; a normative standard is something like being as tall as one ought to be in order to play the leading role in some movie.

⁷ Williamson's (2000) "knowledge-first" programme might be seen as an alternative approach in contemporary epistemology. However, it takes knowledge as primitive and so eschews the project of explaining the difference between true belief and knowledge. Moreover, it carries similar normative commitments in that it sets out knowledge as a normative standard for belief, which raises many of the same questions I will discuss here but in a different key. I set this rich programme to one side in what follows.

⁸ Compare McHugh (2012a, 2012b).

What does it take, in general, for a standard to be “normative” in this evaluative rather than prescriptive sense? I don’t really know. In fact, I dislike the idea of so-called evaluative norms, and would prefer to insist on a stricter separation of normative notions from evaluative notions. But I think there is something right about the idea that some ‘ought’ claims advert to a standard of something’s being good or ideal in some respect, without prescribing that thing’s doing anything. Applied to epistemological ‘ought’s then, the idea is to argue that when we make claims about what someone ought to believe, we are adverting to a “normative” standard for belief regarding what would make a belief good or ideal in some particularly epistemological respect, without prescribing that belief. And this fits well with a normative approach to the theory of knowledge which seeks to identify the “normative” standard, whose satisfaction makes the difference between cases of true belief and cases of knowledge.

Moreover, this idea offers a different route for responding to the ‘ought’ implies ‘can’ objection. For we often deploy these so-called “normative” standards in evaluating things without implying that those things are within the purview of voluntary control, free choice, autonomy. Human intestines ought to maintain a range between 4.0-7.0 Ph. It’d be silly to object to this claim by arguing that human intestines don’t freely choose their Ph values or can’t manifest autonomous agency in maintaining them; and, while humans may have some indirect control over the Ph value of their digestive system through their dietary choices, the ought-claim above does not usually seem like an application of a normative standard meant to indirectly guide dietary choice. Similarly, maybe there are kinds of beliefs human minds ought to maintain (given some body of evidence), and we should understand this claim as adverting to an evaluative rather than prescriptive epistemological norm. If that’s right, then maybe we could respond to the objection to norm governance of belief by maintaining that the relevant ‘ought’s don’t imply ‘can’, in the relevant sense, and so there is no problem with accepting the idea that belief is not voluntary, freely chosen, or autonomous but nonetheless subject to epistemological norms.

Despite not liking the notion of a “merely evaluative” norm, I like this response to the objection so much that I’ve published a paper defending a variant of it (Chrisman 2008). What I particularly like is the fact that, because beliefs are states, claims about what someone ought to believe seem apt to treat as claims about how someone ought to be rather than what someone ought to do. Moreover, if we conceive of knowledge as true belief which also meets some standard of being what someone ought (epistemologically) to believe, I think it seems right to think of this ought-claim not as prescribing belief formation but as evaluating a belief already formed for how well it meets some standard.

I’m sensitive, however, to the worry that belief is importantly different from the Ph values of our intestines. So-called evaluative norms applied to the latter are purely physiological and can be identified by considering the biological function of intestines and determining what Ph values promote that function. Can we do something similar to identify evaluative norms for belief? Maybe, but it is much less clear what function beliefs perform and so more difficult to derive claims about what one ought to believe from a view about the function of belief. Sure, we can say beliefs are for storing

information that may be useful in navigating the world in pursuit of the satisfaction of our desires. But I doubt that all beliefs are for performing that function (moral beliefs, religious beliefs, philosophical beliefs and pure mathematical beliefs all seem like salient counterexamples). And, anyway, epistemological evaluations of belief are not sensitive to whether or not maintaining them promotes desire satisfaction, they're sensitive to considerations about what epistemological reasons there are to believe something, which are the kinds of reasons that relate to knowledge not practical success in satisfying desires.

Partly because of this contrast, in my earlier paper, I argued that epistemological 'ought-to-believe's were mainly interesting for how they related to prescriptive claims about what a believer's epistemic community, which might include him/herself but which also includes other people, ought to do about the believer having the beliefs he/she ought to have.⁹ In this way of thinking of things, evaluative 'ought's about belief get replaced with prescriptive 'ought's about inquiry and instruction. I still think there is a kernel of truth in this idea, and it figures in the more social sort of approach to the theory of knowledge that I want to motivate below. As far as it has been developed so far, however, I now worry that the idea that claims about what someone ought to believe advert to norms about how inquirers and instructors ought to behave risks collapse into to some version of the indirect control response I want to discuss next. For it sounds as if 'ought-to-believe's inherit their normativity from their indirect connection to various prescriptive norms involved in the governance of actions of inquiry and instruction.

4. Indirect Control

The next response I'll consider is an expansion and refinement of the supposedly traditional picture of our mental lives. The idea is to argue that, as long as paradigmatic belief is subject to indirect voluntary control, even if this doesn't often come in the form of deliberation, that's enough to secure the kind of autonomous agency over belief needed to answer the objection. In more detail, someone pursuing this line of thought might argue that many states are like belief in that one often just finds oneself in them: e.g. states like being fit, being rich. But we should recognize an "indirect" kind of control we exercise over such states. We understand more or less how to get into one of these states and how to get out of these states if that's what we take ourself to have reason to do. (Of course, doing so is not always easy, and we are may be subject to weakness of will in our attempts to do so.) If it's right that there are states we more or less know how to get into, it shows that being in a state like one of these is the sort of thing one can decide to do, or do voluntarily, or exercise one's will in doing. In sum, the response is that being fit and being rich are not literally things one *does* in some active way, but they are possible products of things one does in response to (perceived) reason for doing so and so objects of *indirect* voluntary control.¹⁰

Does this provide the resources needed to rebut the 'ought' implies 'can' objection to holding individual believers to normative standards? I think we should recognize that

⁹ Compare Kauppinen (2018) for a similar idea linking the genuine normativity of standards for belief to the sorts of accountability appropriate in the believer's epistemic community for those who violate the standard.

¹⁰ See Meylan (2015, 2017) for a powerful defence of this idea.

there are cases where one does not literally decide to believe but one does decide to actively do something in order to make up one's mind about something. I already mentioned the seemingly special case of belief formed through explicit deliberation about reasons for and against some proposition. But maybe we should also count beliefs as *indirectly voluntary* when they are formed via other activities one engages in to find things out: e.g. investigating the source of some sound, looking things up on Wikipedia, listening to the news, reading a book, etc. As such, it might begin to look like we have something more than just an exception to prove the rule. If these actions are all *indirect* ways to exercise control over one's beliefs, then perhaps we do have enough "autonomous doxastic activity" on the scene to make sense of subjecting believers to normative standards in evaluating whether or not they have knowledge.

Like I said above, I find something attractive in each of the responses to the 'ought' implies 'can' challenge I'll discuss here, but I'm not completely satisfied with any of them. Regarding this response, I have two objections. My first objection is that there is an important difference between belief and the other states given as examples of indirect voluntary control. The normative claims "I ought to be fit" or "I ought to be rich" can guide activity which even if not always easy to execute has the predictable result of one's being in the state. Belief, it seems, doesn't work this way. At least, it is not usual to think "I ought to believe p" and for this normative thought to guide activity resulting in forming the belief p.¹¹ That's sort of how Pascal suggested one might come to believe in God, but this was for explicitly practical reasons, making the 'ought' one of prudence rather than epistemology. Normally, things go the other way around when belief is subjected to epistemic norms. One believes p and then can reflect on whether that's a proposition one ought (epistemically) to believe, perhaps adducing considerations that explain why this is a belief one ought to maintain. Similarly, except in pathological cases like compulsive obsession, if one thinks that one ought (epistemically) not to believe p, that is not anything like the beginning of a plan to get out of the state of believing p but rather the event whereby one ceases to believe. So even though there are states over which we exercise a kind of indirect voluntary control, suggesting that it makes sense to say that one ought to be in them, belief doesn't seem to interact in the same way with voluntary control in response to norm guidance.

One might object that I've misconstrued the way norms guide activity that results in belief. Of course one doesn't usually think, for any specific proposition p, "I ought to believe p" and then set about to acquire that belief. But don't we often think things of the form "I wonder what I ought to believe about x" and then set about investigating x, which we know will result in various beliefs and which we hope to be the beliefs we ought to have about x?

Yes, that's certainly something we sometimes do. But I'd suggest it's not – and this is my second objection to the indirect-control response to the ought-implies-can worry – how most beliefs are formed. I don't really know how to quantify "most" here, but investigation into a topic x with the aim of forming the beliefs about x that one ought to have strikes me as only slightly more common than explicit deliberation about reasons

¹¹ Compare Nolfi (2014) for a more realistic account of how normative evaluation of belief might sometimes affect one's stance towards this belief and the formation of other beliefs. I regard this as a special case and by no means the usual way that belief is formed.

for and against some proposition. When you see someone walking by or you're told that a package will arrive tomorrow or you understand that racism is enjoying political mainstreaming, these are ways of coming to believe something without deliberation *or* investigation.¹² I think epistemologists can and should seek to understand what's going on in cases of explicit deliberation and investigation; and it seems to me that there is an interesting form of indirect voluntary control over our beliefs in such cases. But my present suggestion is that other cases which do not involve active deliberation and investigation are also important and central cases of belief formation, and we should demand that a theory of knowledge and doxastic autonomy treat them as such.¹³

5. Active Beliefs

The final response to the 'ought' implies 'can' objection that I want to consider is to argue that there is a sense in which believing itself is active after all. And once we recognize the distinctive activity in believing, we will be in a position to understand a kind of autonomous activity that is importantly different from autonomous action but nevertheless sufficient to underwrite holding believers to normative standards.

Above I pointed out that belief is typically conceived as a mental *state*, a way one *is* rather than something one *does*. So how could belief itself nevertheless be active? The basic idea in this response is to argue that to understand what it is to believe *p*, we have to appreciate the activities of mind wrapped up with *maintaining* a belief.¹⁴ One doesn't

¹² Compare Nolfi (2018: 67-70) for a related argument against the idea that most doxastic agency is exercised in manipulation of one's environment, e.g. through exposing oneself to various pieces of evidence or putting into place environmental safeguards against believing wrongly.

¹³ Objection: Perhaps beliefs, which are indirectly controlled through deliberation or investigation, are not statistically common but still paradigmatic or normal. After all, for any proposition *p* one believes, can't one go about investigating whether *p* is true or deliberate about whether *p* is true? If yes, we might think that the possibility of indirect control via deliberation or investigation into whether *p* is true is partly constitutive of a mental state's being the belief that *p*, even if most beliefs are not ones actually formed through deliberation or investigation. Response: I'm not so sure one can, even in principle, subject everything one believes to rational reflection characteristic of deliberation and investigation. Maybe some beliefs are ineffable. And even if one can, in principle, deliberate or investigate whether *p* for any *p* one believes, I suspect one will often need social input to determine what one believes. (More on this below, when I consider a similar idea that one can in principle rationally scrutinize any belief one holds.) Moreover, this objection trades on difficult issues about how to understand what constitutes a mental state's being a belief. Since we can't open up the mind, identify the beliefs, and investigate their nature, I suspect this question is ultimately a task for conceptual engineering. That is to say that we face a question about what conception of belief is best suited for our theoretical purposes. I'm content to recognize a diversity of theoretical purposes for the concept of belief, and maybe the conception of belief advanced by this objection is well suited for some interesting theoretical purpose. However, I'm also inclined to think that this conception is not well suited for the purpose of understanding knowledge in its full generality. Meylan (2017: 14-15, 19-21) argues that we exercise indirect agency far more often than it seems and that we can foresee many doxastic consequences of our actions, and I don't mean to deny this. It's just that many of what I see as paradigmatic cases of knowledge (e.g. via perception or testimony or gradual settling into a view that just seems to make sense of a lot of things) don't seem to me to be like that. Compare Owens (2000: 87) for a different sort of critique of the indirect agency response to the objection.

¹⁴ See Boyle (2011) for defences of the idea that belief is an active state. I address Boyle's arguments more directly in Chrisman (2018). See also Hunter (2018) for criticism of Boyle and criticism of the view

typically (ever?) pick up and hold a single belief isolated from other beliefs, impressions, and actions. When one believes *p*, one is defeasibly disposed to immediately form beliefs about the obvious consequences of *p* if the question of their truth is relevant or salient. One is defeasibly disposed, if one has any cause to reflect on it, to feel convinced of *p* and take there to be evidence for this conviction; one is also defeasibly disposed to distrust those who assert not-*p* and to be surprised if one has an experience as if not-*p*. One is also defeasibly disposed to perform various cognitive actions such as to treat *p* as evidence in deliberations and investigations into other things, rely on the truth of *p* in various forms of practical reasoning, and reflect on the grounds for believing *p* if one discovers this belief is inconsistent with something else one believes.¹⁵

As indicated, all of those dispositions are defeasible, and there are surely others than the ones I just mentioned. But the key point is that our grip on paradigmatic cases of believing *p* seems to depend crucially on appreciation of some complex disposition to engage in various mental and nonmental activities involved in maintaining the belief *p*. So, even if we want to insist (in full metaphysical precision) that belief is a state and so not a paradigm activity, we should recognize that it is a state we understand only via a tacit understanding of a complex disposition to engage in intimately related activities. And, according to this second response to the objection, that's enough to make belief active by proxy in a way that most states are not. As long as maintaining a belief is something one does actively, then it can make perfectly good sense to hold people to normative standards about what they ought to believe. For believing as one ought to is the same thing as maintaining the beliefs one ought to maintain.

As before, I think this response to the 'ought' implies 'can' objection is instructive, but I am not fully satisfied. To begin to explain why, let me point out that even if we count belief itself as active "by proxy" of the activity of maintaining belief, that doesn't show that we decide to believe, voluntarily believe, or exercise our wills in believing, or anything similar, which is what the 'ought' implies 'can' objection appears to demand.

I outline in that paper. Here I'm considering what I take to be an attractive descendent of Boyle's idea. It's different from the similar ideas floated in Chrisman (2018) that doxastic agency is located most centrally in the activity of maintaining a whole system of beliefs and Nolfi (2018) that doxastic agency is located most centrally in the cognitive processes manifesting our belief-regulating dispositions. In both cases, these views focus on whole belief systems rather than individual beliefs, whereas Boyle focusses on individual beliefs. The picture I develop below is an extension into the social sphere of the idea of locating doxastic agency at the level of systems of belief, which I now view as inherently social in beings like us. Hieronymi (2009) posits a *sui generis* kind of activeness involved with continually settling a question for oneself that she thinks belief manifests. One way to make of this idea is as the view that continually settling a question whether *p* for oneself is a way of engaging in the activity of maintaining a belief. If that's what Hieronymi has in mind, it is a version of the idea I discuss in the text above. In fn 17 below I consider a different way of making sense of that idea based in Hieronymi's idea of "evaluative control"; see also Setiya (2008) for different critical discussion of Hieronymi's attempt to make sense of the activeness of belief.

¹⁵ Schwitzgebel (2002) develops an account of belief as a complex dispositional state. He argues that there is a stereotype of dispositions associated with various beliefs, and one can manifest the stereotypical dispositions involved in believing *p* to a greater and lesser degree.

To see this, notice that there are lots of states we understand only (or at least primarily) by reference to correlated dispositions to activities.¹⁶ What is it to be frightened, for instance, except to be disposed to flee, perspire and have a faster heart rate, form impressions of danger, etc.? Or what is it to be awake except to be disposed to consciously experience certain kinds of things, engage in certain kinds of brain activity, etc.? I don't know the right account of being frightened or being awake. But my point in mentioning these ideas is that, even if it's true that we can make sense of these states only via an appreciation of a complex dispositions to engage in the activities intimately related to them, that doesn't show that we typically decide when to be in one of these states or that being in one of those states is voluntary in some sense or that being in one of these states involves some kind of exercise of our will. We can, of course, voluntarily choose to do things that we know to eventuate in being in these states (e.g. watch the horror film to induce fright, drink the strong cup of coffee to maintain wakefulness). Modeling doxastic autonomy in this fashion, however, is back to the indirect control response to the objection explored in the previous subsection.

More to the point, however, even if believing *p* is to be understood by essential reference to the activities wrapped up in maintaining the belief *p*, a proponent of the picture of our mental lives dominant in contemporary cognitive psychology and behavioral economics is still going to argue that a whole lot of the beliefs that we maintain are automatic, intuitive, and/or the products of implicit biases and ideological framing. Perhaps, they might grant, when you see that there are words on this page, you're engaged in a lot of cognitive and other activity in maintaining the belief that there are words on this page. But still, taken as a whole complex of conditional dispositions, it's not up to you whether or not to engage in this activity, so why can we hold you to a normative standard regarding this activity? Why think this activity manifests your autonomous agency? Similarly, maybe my maintaining the belief that the currency of Mexico is the peso involves activating a complex disposition to various activities, some of which are clear cases of voluntary action, such as exchanging money for pesos when I cross the border into Mexico. But still, one might reasonably insist that it's not up to me whether to maintain this belief, so why do I get held to a normative standard regarding it?

More abstractly, the worry about the active belief response to the objection is that, although some of the activities dispositionally connected to believing *p* may be voluntary or things we choose to do, that doesn't really address the objection, unless one can choose to maintain or can voluntarily maintain or can exercise one's will in the very activity of maintaining a belief. And, so far at least, it doesn't seem that one can do this in very many of the paradigm cases of belief.¹⁷

¹⁶ See Hunter (2018) which makes a similar point.

¹⁷ Some might want to argue that we are nonetheless active in maintaining a belief in a way that distinguishes belief from other mental states such as headaches. For example, Hieronymi (2006) argues that we enjoy "evaluative control" over our beliefs insofar as our maintaining them is responsive to "constitutive reasons" which bear on the truth of their content, which she distinguishes from the "managerial control" involved in responding to "extrinsic reasons" for having various beliefs, i.e. acting so as to come to have various beliefs because it will be good to have them. Since such evaluative control is not a kind of control we have over headaches, does this mark out a distinctive form of mental activity? Sure, but I'm less convinced that it marks out a kind of agency, as most reasons-responsive belief still

In response, one might object that I've mischaracterized the activity of maintaining belief or at least the bit of it where voluntary control gets into the picture. For there seems to be a potential mental activity one can voluntarily engage in with respect to any belief. One can *notice* what one believes about some topic, and then *subject that belief to further scrutiny*. Even Descartes, arguably, wasn't suggesting in the *Meditations* that we can start with a blank slate and then go about adding beliefs to it at will. Rather he is more plausibly read as suggesting that we enjoy the freedom to take a reflective stance towards each of our beliefs, gathering further evidence for and against it, and deliberating about whether it is a belief we ought to have. If we count as autonomous only those beliefs for which one has actually engaged in such reflective activity, then the sphere of doxastic autonomy will be too small. If, however, maintaining belief somehow involves a stance towards one's beliefs whereby this kind of reflective activity is always possible (even if rarely actual), we may have hit on a conception of the activity of maintaining a belief that extends the sphere of doxastic autonomy wide enough to rebut the 'ought' implies 'can' objection.¹⁸

I think this is the best version of the active belief response to the objection, and I am very sympathetic. However, I suspect Descartes was wrong if he thought it generally humanly possible to sit alone by the proverbial fireplace and notice from a third-personal perspective what beliefs one has about any topic, in order to then subject them to further scrutiny, identifying the bases of these beliefs and considering the support they provide. This is a more general challenge for the idea that we are autonomous in maintaining our beliefs because we always *can* subject anything we believe to rational reflection and investigation. Even identifying *what* one believes about many topics requires intercourse with other people and is sometimes very difficult, possibly impossible, to do accurately.¹⁹ After all, if most belief is automatic, intuitive, biased, and/or ideologically framed, there's no guarantee one can just, so to speak, ask oneself what one believes and get a straightforward answer. Moreover, many of the concepts we deploy in believing are externalist, in the sense that their application conditions depend on connections to things in the world including the thoughts and actions of other people.

seems to me to be intuitive, automatic, and unchosen. Moreover, I suspect that there are mental states such as being frightened or being bored that are often responsive to constitutive reasons without manifesting the sort of agency presupposed by holding someone to genuine normative standards. When we say that someone ought not be frightened by the thunderclap, we don't blame them for being frightened like we blame someone for not believing the obvious consequences of what they already believe. Compare McCormick (2018) who allows that evidence-responsiveness might not be enough to secure doxastic agency but goes on to argue that we can, in some intuitive cases, believe in response to non-evidential reasons which is characteristic of paradigmatic exercises of agency and which would distinguish believing from being frightened or being bored.

¹⁸ Keeling (forthcoming) suggests that one can, at least in principle, determine *why* one believes what one believes for any belief one holds. She contends that, even if we do not exercise choice or control in forming a belief, we can always ask for anything we believe "Why (do I) believe p?" And she thinks taking up the deliberative stance and answering "Because q" is a matter of choosing the basis for our continuing to believe p. This is different from the diagnostic stance, whereby one might try to explain what caused one to have the belief in the first place. And the answer from the deliberative stance is something one can autonomously do (choose the basis) which is also causally efficacious (insofar as one does indeed continue to believe p).

¹⁹ Compare Bem (1970), Gopnik and Meltzoff (1994), Lycan (2008), Dretske (2004), and Carruthers (2009) for various arguments against the idea that we can generally introspect what we believe.

And even more importantly, when it comes to scrutinizing a belief, I think one usually needs to leave the armchair. Of course, one needs to leave the arm chair to gather further evidence about empirical facts, but the world doesn't present its facts to the lonely explorer unvarnished by the natural and social influences on our belief forming processes. So one needs to leave the armchair and link minds with other believers, in order to go out and investigate the world.

This is part of why I believe we need a more socialized approach to the theory of knowledge and understanding the kind of doxastic autonomy enjoyed in certain kinds of epistemic communities. I turn to that in the next section. I want to take a step back to the theory of knowledge approached via a normative standard, but I will suggest this approach is best pursued not in terms of a standard individual beliefs must meet but rather in terms of the role doxastic standards play in generating a kind of autonomy available only through our sociality.

6. Doxastic Autonomy as Social a la Rousseau

So far I have been exploring an objection to an approach to the theory of knowledge which assumes that the key question in differentiating true belief from knowledge is one about the normative standard believers must meet in order for one of their beliefs to count as knowledge. The objection, in short, is that holding believers to normative standards assumes that believers exercise autonomous agency in holding their beliefs, but this assumption is wrong. Or at least it is mostly wrong: when we look at any individual belief, it will usually be the product of automatic, intuitive, biased, and/or ideologically-influenced cognitive processes; and that's enough to undermine any assumptions of doxastic autonomy.

Although I think the responses considered above have some merit, I want to suggest that a full response requires reconceptualizing doxastic autonomy in a more social key. I propose we consider moving away from a focus on individuals and what they can or cannot freely believe to a focus on communities and the way norm-governance within the context of our sociality facilitates different kinds of autonomy. My idea is to base a fuller understanding of doxastic autonomy in the sort of epistemic sociality involved in our collective enforcement of particular sorts of rules. This is based on an analogous understanding of practical autonomy that emerged in Rousseau's (1762) response to Hobbes' (1651) account of the genealogy of political morality. So next I'll give a thumbnail sketch (with no pretensions of exegetical preciseness) of that important moment in the history of philosophy.

Famously, Hobbes imagined pre-moral humanoids to be free in the sense that they are not subject to any moral rules which could be enforced with legitimacy. He then tried to explain how such pre-moral humanoids could have moved from the state of nature into a society where at least someone was in a position to legitimately enforce something recognizable as moral norms. The basic strategy in his explanation was to identify the "natural" freedoms a characteristic individual pre-moral human would be rational to give up in exchange for the security which comes with submission to the social enforcement of a set of rules. And on this basis, he suggested that we should think of

ourselves (contemporary and fully social humans) as having autonomously albeit implicitly accepted the legitimacy of social enforcement of those rules, which do limit our freedom.

In response, Rousseau granted that pre-moral humanoids would possess a kind of “natural” freedom which we contemporary and fully social humans lack, but he argued that such freedom is different from genuine autonomy. This is because the activities we fully socialized humans characteristically engage in, and want to engage in autonomously, can only be conceptualized in ways that already presuppose commitment to collective projects which are partly constituted by (tacit) acceptance of norm governance. For example, in Rousseau’s view, genuinely autonomous humans are ones who can do things like get married, play a football game, purchase a vacation home, sing in a choir, say #metoo, convict a criminal, fight climate change – activities all of which require cooperation with other people recognized as moral agents to whom one gives the authority to enforce various norms. In this way, Rousseau conceived of the kind of freedom characteristic of autonomy as a possible result of rather than precondition on our sociality as humans subject to the authority of moral rules.

There is significant discussion in the history of political philosophy about Hobbes’ and Rousseau’s contrasting approaches to political autonomy, and my intention isn’t to contribute to this discussion here. Rather I mention this contrast here in the context of thinking about the kind of autonomy presupposed in holding believers to normative standards because I think it may provide a new perspective on what that kind of autonomy looks like. And this perspective might shed new light on the role of normative governance of belief in the theory of knowledge.

More specifically, in thinking about doxastic autonomy, I think it might prove useful to reflect on the kinds of social arrangements that make for more and less autonomous communities of believers. This is in parallel to (though obviously also deeply intertwined with) ideas we might have about the kinds of social arrangements that make for more and less autonomous communities of moral agents. For instance, within a society of people structured wholly as masters and slaves, the slaves almost completely lack practical autonomy. Similarly, we might think that wherever some people in a community are wholly dominated in thought through indoctrination, these people almost completely lack doxastic autonomy.²⁰ On the other side, however, in a society where there are no enforceable norms for how people treat each other, complete anarchy reigns and we shouldn’t regard this kind of society as fostering practical autonomy. Similarly, if there are no enforceable norms for how people treat each other intellectually, something akin to cognitive anarchy might reign and we wouldn’t then regard this community as mastering enough common understanding of reality to foster doxastic autonomy.

That quick sketch is highly unrealistic given strong evolutionary and social forces that seem to prevent extreme domination or complete anarchy, in both the practical and

²⁰ This is related to the republican conception of freedom as non-domination (see especially Pettit 1997: ch. 2 and Pettit 2012: ch. 1); however, I go on to suggest that one also lacks doxastic autonomy in an intellectual anarchy. Compare also Medina’s (2013) appeal to the value of epistemic friction for resisting the threats of epistemic domination.

doxastic realms. I hope, however, that the sketch can help to frame an investigation into the elements of a social arrangement that affect how much practical and doxastic autonomy we think there is in a group of people. In order to understand this kind of autonomy, we need to find the middle ground between domination and anarchy, in both the practical and doxastic realms, both of which I'm now conceiving of as essentially social. And this middle ground lies, I want to suggest, in the social arrangements that allow for collective pursuit of common practical and intellectual aims, and these social arrangements are ones constituted by (implicit) acceptance of norm governance.

What are the common practical and intellectual aims? In its most abstract formulation we might conceive of these aims as *living well together* and *understanding the world* (again, obviously, deeply intertwined). Maybe that is too abstract to be anything but a placeholder for a theory, but we can begin to make progress on that theory in the case of collective intellectual aims by noticing that humans (essentially?) have diverse perspectives on the world and different capacities to expertise; and we improve our understanding of the world by sharing our diverse perspectives and integrating our respective expertise. To put it simplistically, you see things from there, I see things from here, we collectively understand what we see by finding a way to share each other's perspectives. You have the capacity to learn about x, I have the capacity to learn about y, we both better understand the world when we each go out and learn about one of these things and then find a way to integrate what we have learned in a common understanding of some aspect of the world. Of course, it's not only or even mainly about you and me. Our epistemic lives involve many overlapping communities of people who share perspectives and integrate expertise as part of improving collective understanding of the world.²¹

I also want to suggest that this sharing of diverse perspectives and integrating respective expertise depends on mutual and collective albeit tacit submission to epistemic norms. In trusting others in one's epistemic community, we accept the legitimacy of various epistemic norms. To share perspectives requires more than mere combination, we must apply standards for addressing tensions or incompatibilities in the diverse perspectives. To integrate differing areas of expertise, we must apply standards for organizing various bodies of knowledge along various explanatory and justificatory dimensions. These standards are epistemological norms.

I'm not claiming that there's one universal set of norms with absolute epistemic authority. Epistemic communities might be usefully evaluated as more and less autonomous depending on how well the norms they mutually (though tacitly) accept promote collective understanding of the world. But there could be different parts or aspects of the world whose understanding is better or worse promoted by different sets of epistemic norms, and there could be different paths to equally good understanding of the world. However, as we move towards the extremes of intellectual domination and

²¹ This may wrongly suggest I think everyone in a community has to have the aim of understanding the same part of the world and that the strategies contributions to this aim are always consensual and cooperative. The overlapping nature of various diverse epistemic communities already suggests that there may be a common aim only in some very abstract sense. Moreover, there are surely different competing strategies for pursuing understanding of any part of the world. Compare Weisberg and Muldoon (2009) on research "followers" and research "mavericks".

intellectual anarchy in imagining different possible epistemic communities, it is less and less plausible that there are mutually accepted epistemological norms in these communities. And the lack of such epistemological norm governance threatens even the most basic or piecemeal collective understanding of the world.

On this kind of picture, knowledge is spread out in complicated ways across communities of believers whereby individuals' knowledge-conferring justification for belief often depends on things understood only by other people. And while it might make sense to think of how much of reality particular individuals understand, the degree to which reality is understood and the process of improving that understanding is a collective matter.²²

What does this tell us about the kind of doxastic autonomy presupposed by holding believers to normative epistemic standards? My suggestion is that we shouldn't think of such norms as presupposing some conceptually antecedent sort of capacity to decide to believe or voluntarily choose our beliefs or freely exercise our will in an individual belief. Rather, following Rousseau's treatment of practical autonomy as emerging from self-submission to normative governance that constitutes collective action, as contrasted with natural freedom, I want to suggest that believers are autonomous (to the degree that they are) in virtue of the way their beliefs contribute to a collective epistemic project of understanding of the world. Participation in this project is constituted by implicit acceptance of the authority of some mutually enforceable epistemological norms. The normative epistemological standards applying to belief on this picture are not prescriptions that could or should guide a believer's belief-formation. This is a way of agreeing with an aspect of the "evaluative"-norms response above to the 'ought' implies 'can' objection. However, I think it needs to be taken further into a social conception of the nature and role of holding beliefs to (evaluative) normative standards.

We may not (usually) get to choose what we believe, but we do get to choose whether to testify to others, contribute what we believe to some group's overall understanding of some issue, check our beliefs against technological-cum-social repositories of collective knowledge, and appeal to others in appreciating the ideologies and biases influencing belief formation and information processing in ourselves and the institutions in which we have power. Those are the sorts of autonomous activities that don't even make sense from the individualistic perspective but which seem to me to be quite central to what it is to be someone who knows something – at least in most of the cases where it might make sense to ask whether someone believes as they ought to believe. This is a way of agreeing with aspects of the indirect-control and active-belief responses above to the 'ought' implies 'can' objection, but I'm suggesting that too they need to be taken further in the social conception of the relevant actions and activities intimately connected to belief.

Above I expressed a worry about conceiving of claims about what someone ought to believe as indirectly about what they ought to do to get themselves to believe something. The worry was, roughly, that the relevant ought-to-believe claims can guide the believer's activity only in cases where one takes something like an alienated or

²² Millgram (2015: ch. 2) presents an especially pessimistic account of this. See Nguyen (2018) for a more hopeful response.

therapeutic stance towards one's own belief. In the normal case, one's belief that p is antecedent to one's belief that p is what one ought to believe. But the situation is obviously different with claims about what someone else ought to believe and claims about what "we" ought to believe. A professor can think that her students ought to believe that Kant wrote the *First Critique*, and there's nothing mysterious or pathological about that thought guiding her actions of instruction. A political commentator can think that we (members of her epistemic community) ought to believe that racism is enjoying political mainstreaming, and there's nothing mysterious or pathological about that thought guiding her actions of investigative reporting.

Rousseau argued that it was through tacit agreement with a social contract that humans free themselves from our animalistic passions and urges and thereby achieve the kind of autonomy needed to pursue the sorts of robust projects characteristic of a good human life. These projects depend on our ability to pursue the common good, which in turn depends on our ability to understand moral norms as rules we give unto ourselves. By extending this conception of autonomy and norm-governance from morality to epistemology, I think we can make sense of a different – and I submit more interesting – sphere of doxastic autonomy from the ones often discussed in debates about doxastic voluntarism and normative epistemology. It is tacit agreement with a social epistemic contract that humans free themselves from the limitation of knowledge to things in their immediate environment and thereby achieve the kind of intellectual autonomy needed to pursue the intellectual projects characteristic of a good human life. Just like our practical projects (and as an integral part of them) these epistemic projects depend on our ability to pursue common knowledge and transmit it culturally, which in turn depends on our ability to understand epistemic norms as rules given unto ourselves.

7. Conclusion: Towards Social Epistemological Approaches to the Theory of Knowledge

At the beginning of this paper, I mentioned the methodology of contrasting cases that has dominated the theory of knowledge, in both its nonnormative and normative guises. By extending the Rousseauian approach to political autonomy to intellectual autonomy, I think we're in a position to start thinking about other methodologies. In some fairly generic and so not terribly informative way, we might begin with the idea that knowledge is true belief one epistemologically ought to have. Then the more social understanding of knowledge encouraged by the conception of doxastic autonomy just outlined would encourage us to think about this 'ought' in terms of what one can contribute to the collective project of understanding the world.²³

That is to say, broadly speaking, what one ought *epistemologically* to believe should be defined in terms of what norms of belief adherence to which will improve our collective understanding of the world. This is the proposed normative difference between mere true belief and knowledge. How do we evaluate which ways of believing will improve

²³ Compare Craig (1990) and Rosenberg (2002) for similar attempts to reconceive the methodology for developing a theory of knowledge.

our collective understanding of the world? Contrasting cases may help, but I'd suggest that a more illuminating methodology will be to think about the sorts of belief we want to encourage in fellow members of our epistemic community as part of sharing diverse perspectives and integrating different areas of expertise. This might include various ways of believing, and different ways might be appropriate for understanding different aspects of reality or within different epistemic communities depending on their histories and structures. But some of the areas I think we should think about more are testimony, memory, extended information systems, just treatment of other believers.

More specifically, if we're investigating the epistemological standards believers should be held to in thinking about whether they count as knowers, we should think about the role such knowledge plays in an economy of information within a community of believers. Framed this way, we might come up with claims like the following: One ought to believe in ways one should stand behind in testimony.²⁴ One ought to believe in ways that cohere not only with reflectively accessible knowledge but also technologically and socially accessible knowledge.²⁵ One ought to believe in ways that avoids epistemic injustice, i.e. in a way that meets the social obligation not to discount the testimony of others because of prejudice against the type of person they are.²⁶ One ought to believe in ways that acknowledge fallibility in the face of implicit biases. One ought to believe in ways that are self-aware about the unavoidable ideological framing of various issues and the challenges this poses for sharing perspectives.

Do these 'ought's imply the 'can' of autonomous activity? The more social approach to the theory of knowledge I am exploring here suggests that the degree to which we hold each other accountable to such social rules (and the kind of sanction/reward structure attaching to that accountability) is a key part of what makes some communities of believers more autonomous and others less. On this picture, particular belief lies outside the scope of individual choice, and the activity of maintaining a belief may be mostly automatic, intuitive and even driven by implicit biases and ideologies. But we should think of that activity as embedded in and structured by social rules constituting epistemic communities. And when these social rules promote mutual understanding of the world, they are the sorts of norms, conformity to which fosters the very kind of doxastic autonomy presupposed in saying that someone (epistemically) ought to believe something.

Bibliography

- Alston, William (1988): "The Deontological Conception of Epistemic Justification", *Philosophical Perspectives* 2, 257–299.
- Anderson, Elizabeth (2012): "Epistemic Justice as a Virtue of Social Institutions", *Social Epistemology* 26(2), 631–73.
- Bem, Daryl (1970): *Beliefs, Attitudes, and Human Affairs*, Belmont, CA: Brooks/Cole.
- Boyle, Matthew (2011). "'Making up Your Mind' and the Activity of Reason", *Philosophers' Imprint* 11, 1–24.

²⁴ Compare Goldberg (2010, 2018).

²⁵ Compare Goldman (1999).

²⁶ Compare Fricker (1997) and Anderson (2012).

- Carruthers, Peter (2009): “How We Know Our Own Minds: The Relationship between Mindreading and Metacognition”, *Behavioral and Brain Sciences* 32, 121–82.
- Chrisman, Matthew (2018): “Epistemic Normativity and Cognitive Agency”, *Noûs* 52 (3), 508-529.
- Chrisman, Matthew (2012): “The Normative Evaluation of Belief and The Aspectual Classification of Belief and Knowledge Attributions”, *Journal of Philosophy* 109 (10), 588-612.
- Chrisman, Matthew (2008): “Ought to Believe”, *Journal of Philosophy* 105(7), 346-370.
- Craig, Edward (1990): *Knowledge and the State of Nature: An Essay in Conceptual Synthesis*, Oxford: Oxford University Press.
- de Sousa, Ronald (1971): “How to Give a Piece of Your Mind: Or, The Logic of Belief and Assent”, *Review of Metaphysics* 25(1), 52-79.
- Dretske, Fred (2004): “Knowing What You Think vs. Knowing That You Think It”, in *The Externalist Challenge*, ed. by Schantz, Berlin: Walter de Gruyter, 389–99.
- Evans, Gareth (1982): *The Varieties of Reference*, Oxford: Oxford University Press.
- Fricker, Miranda (2007): *Epistemic Injustice: Power and the Ethics of Knowing*. Oxford: Oxford University Press.
- Goldberg, Sanford (2010): *Relying on Others: An Essay in Epistemology*, Oxford: Oxford University Press.
- Goldberg, Sanford (2018): *To the Best of Our Knowledge: Social Expectations and Epistemic Normativity*, Oxford: Oxford University Press.
- Goldman, Alvin (1999): *Knowledge in a Social World*. Oxford: Oxford University Press.
- Gopnik, Alison and Meltzoff, Andrew (1994): “Minds, Bodies and Persons: Young Children’s Understanding of the Self and Others as Reflected in Imitation and ‘Theory of Mind’ Research”, in *Self-Awareness in Animals and Humans*, ed. by Parker *et. al.*, New York: Cambridge University Press, 166–86.
- Hieronymi, Pamela (2006): “Controlling attitudes”, *Pacific Philosophical Quarterly* 87(1), 45-74.
- Hieronymi, Pamela (2009): “Believing at Will”, *Canadian Journal of Philosophy*, supplementary volume 35, 135-187.
- Hobbes, Thomas (1651/2016): *Leviathan*, ed. by Missner, New York: Routledge.
- Hume, David (1739-1740/2002): *Treatise of Human Nature*, ed. by Norton/Norton Oxford: Oxford University Press.
- Hunter, David (2018): “The metaphysics of responsible believing”, *Manuscrito*, 41, 1-31.
- Kauppinen, Antti (2018): “Epistemic Norms and Epistemic Accountability”, *Philosophers’ Imprint* 18, 1-16.
- Keeling, Sophie (forthcoming): “The Transparency Method and Knowing our Reasons”, *Analysis*.
- Lycan, William (2008): “Phenomenal Intentionalities”, *American Philosophical Quarterly* 45:3, 233–52.
- Mandelbaum, Eric (2016): “Attitude, Inference, Association: On the Propositional Structure of Implicit Bias”, *Noûs* 50(3), 629-658.
- McCormick, Miriam (2018): “Responding to Skepticism about Doxastic Agency”, *Erkenntnis*, 83(4), 627-645.

- McCormick, Miriam (2015): *Believing Against the Evidence: Agency and the Ethics of Belief*, New York: Routledge.
- McCormick, Miriam (2011): "Taking control of Belief", *Philosophical Explorations* 14(2), 169-183.
- McHugh, Connor (2012a): "Epistemic Deontology and Voluntariness", *Erkenntnis* 77(1): 65-94.
- McHugh, Connor (2012b): "The Truth Norm of Belief", *Pacific Philosophical Quarterly* 93(1), 8-30.
- Medina, José (2013): *The Epistemology of Resistance*, Oxford: Oxford University Press.
- Mercier, Hugo and Sperber, Dan (2017): *The Enigma of Reason*, Cambridge, MA: Harvard University Press.
- Meylan, Anne (2015): "The Legitimacy of Intellectual Praise and Blame", *Journal of Philosophical Research* 40, 189-203.
- Meylan, Anne (2017): "The Consequential Conception of Doxastic Responsibility", *Theoria* 83(1), 4-28.
- Millgram, Elijah (2015): *The Great Endarkenment: Philosophy for an Age of Hyperspecialization*, Oxford: Oxford University Press.
- Nolfi, Kate (2014): "Why is Epistemic Evaluation Prescriptive?", *Inquiry: An Interdisciplinary Journal of Philosophy* 57(1), 97-121.
- Nolfi, Kate (2018): "Moral Agency in Believing", *Philosophical Topics* 46(1), 53-74.
- Nguyen, C. Thi (2018): "Expertise and the Fragmentation of Intellectual Autonomy", *Philosophical Inquiries* 6(2), 107-124.
- Owens, David (2000): *Reason without Freedom: The Problem of Epistemic Normativity*, New York: Routledge.
- Pettit, Philip (1997): *Republicanism: A Theory of Freedom and Government*, Oxford: Oxford University Press.
- Pettit, Philip (2012): *On the People's Terms: A Republican Theory and Model of Democracy*, Cambridge: Cambridge University Press.
- Rosenberg, Jay (2002): *Thinking About Knowing*, Oxford: Oxford University Press.
- Rousseau (1762/2002) *The Social Contract*, ed. by Dunn, New Haven: Yale University Press.
- Schwitzgebel, Eric (2002): "A Phenomenal, Dispositional Account of Belief", *Nous* 36(2), 249-275.
- Setiya, Kieran (2008): "Believing at Will", *Midwest Studies in Philosophy* 32, 36-52.
- Stevenson, Leslie (2002): "Six Levels of Mentality", *Philosophical Explorations* 5(2), 105-124.
- Steward, Helen (1997): *The Ontology of Mind*, Oxford: Oxford University Press.
- Tversky, Amos and Kahneman, Daniel (1973): "Availability: A heuristic for judging frequency and probability", *Cognitive Psychology* 5(2), 207-232.
- Velleman, J. David (1985): "Practical reflection", *Philosophical Review* 94(1). 33-61.
- Weisberg, Michael & Muldoon, Ryan (2009): "Epistemic Landscapes and the Division of Cognitive Labor." *Philosophy of Science* 76(2), 225-252.
- Williams, Bernard (1970): "Deciding to Believe, in his *Problems of the Self*, Cambridge: Cambridge University Press.
- Williamson, Timothy (2000): *Knowledge and its Limits*, Oxford: Oxford University Press.