

ARTICLE

Social Epistemic Normativity: The Program

Sanford C. Goldberg

Department of Philosophy, Northwestern University, Evanston, IL, USA

Email: s-goldberg@northwestern.edu

(Received 11 January 2019; accepted 18 November 2019)

Abstract

In this paper I argue that epistemically normative claims regarding what one is *permitted* or *required* to believe (or to refrain from believing) are sometimes true in virtue of what we owe one another as social creatures. I do not here pursue a reduction of these epistemically normative claims to claims asserting one or another (ethical or social) interpersonal obligation, though I highlight some resources for those who would pursue such a reduction.

Keywords: epistemology; social epistemology; norms; obligations; reduction

1. Evaluation, prescription, and normativity in epistemology

For present purposes, an epistemological theory is *evaluative* just in case it formulates evaluative standards by which to assess its targets – cognitive states of belief, credence, or what-have-you – from the epistemological point of view. When applied to these targets, an evaluative theory will imply verdicts regarding the states' epistemic goodness or badness (relative to the evaluative standards of the theory). I will regard an epistemic theory as *normative*, by contrast, just in case it formulates conditions under which S is permitted, required, or forbidden from having such-and-such a doxastic attitude.

While most epistemologists will be happy to acknowledge the evaluative dimension of epistemology, it is a controversial matter whether epistemological theory is, or should be, normative in the sense I intend here. Clarity will be served by bringing out the nature of the controversy. Suppose you think that an evaluative epistemological theory will include the standards of rationality, and so (as applied to particular cases) will yield verdicts about which doxastic states are rational (and which are not). You might then be tempted to conclude that such a theory is already a normative one, insofar as it yields verdicts about which doxastic states are rationally *permissible* and which are rationally *required*. But this is not yet enough to show that the theory is normative in the sense intended here. Consider: a theory of etiquette will include the standards of proper etiquette, and so can be seen as a theory of what is permissible or required, etiquette-wise. But even if something is required etiquette-wise, it may not be incumbent on one to perform the act (or bring about the state of affairs) in question. To be normative in the sense intended here, a theory must be such that, as applied to particular cases, it yields categorical verdicts about what is permitted or required. An epistemological theory is normative, then, just in case it yields categorical verdicts about which doxastic states a subject is permitted or required to have.

© The Author(s), 2020. Published by Cambridge University Press

As I say, it is a controversial matter whether epistemology is normative in this sense. Those who think it is, and so who think that there are doxastic states which we are categorically permitted or required to have, owe an account of what underwrites such normative verdicts (as I will call them). My aim in this paper is to propose a source for epistemically normative claims regarding what we are categorically permitted or required to believe. My contention will be that there are a good many cases in which the underwriting material is social: *we owe it to one another* to conform to the evaluative standards of epistemology.

It will be recognized immediately that this proposal is a variant on a rather traditional approach. I have in mind Clifford's account of the ethics of belief. Arguably, he endorsed a broadly social approach to epistemology's normative verdicts, treating claims about what we ought to believe (= what we are categorically required to believe) as a special kind of ethical requirement. In the cases that I will discuss my analysis will be very much like Clifford's: I will be construing epistemically normative claims about what we are permitted or required to believe as underwritten by claims about what we owe to one another. But my view will differ from Clifford's in at least three significant ways. First, whereas Clifford himself was an evidentialist about the evaluative standards of epistemology, I aim to be neutral on this. Second, whereas Clifford held that the sort of normativity at issue reflected our ethical values, I will be neutral as between whether the basis of epistemic permissions/requirements is exhausted by our ethical values or also include the value of (cooperative) rational agency itself (where this is conceived to be a value independent of our ethical values). Third, whereas Clifford made a claim about the nature of epistemic normativity itself, aiming thereby to characterize the nature of the phenomenon, my aim is less ambitious: I only argue that there are many cases in which what you are categorically required or permitted to believe reflects what you owe others; such a claim is consistent with alternative accounts of the nature of epistemic normativity itself.

The position I wish to defend here can be stated programmatically as follows. Whatever the correct evaluative standards of epistemology are, there are a great many cases in which each of us ought to avoid violating these standards in our doxastic lives *because others are entitled to expect as much from us*. This entitlement to expect one another to live up to the evaluative standards of epistemology is not an *epistemic* entitlement; it is rather the sort of entitlement that we have to expect others to refrain from behaving in ways that undermine our agency or that render them unfit for future cooperation. It reflects our nature as social creatures who systematically and ineliminably rely on one another, where the reliance in question implicates our doxastic states in various ways.

One hope I have with this paper is to identify resources for those who think that epistemic requirements are *essentially* interpersonal in nature. While I do not defend such a strong view here, I offer various resources that might be used by those who seek to defend the stronger thesis.

2. Towards a more social normativity for epistemology

Suppose we have arrived at a correct theory of the evaluative standards for epistemology. We can then use these standards to assess, in a given case, whether a subject's doxastic attitudes¹ conform to those standards. Let us stipulate that doxastic attitudes that conform to these standards are epistemically good and those that do not are epistemically bad. We can then ask: in what sense, if at all, *ought* we conform to epistemic standards when we form our doxastic attitudes? Are we *required* to do so?

¹These include credences. However, in what follows I will no longer mention these.

Those who distinguish the evaluative from the normative have long been aware that not all evaluative standards are such that we are required to conform to them. Above I gave the example of the evaluative standards of etiquette, but there are many others: consider the standards of fashion, dance, effective criminality, landscape architecture, and so forth. In none of these cases is it categorically required of one to conform to these standards. Rather, one conforms to them only insofar as one has an antecedent desire to do well in the domain in question.

Many people think that the evaluative standards of epistemology are different – that the demand that we conform to them is categorical, and that in any case this demand does not rest on anything we happen to want or value.² To see how natural this is, suppose that a good many of your beliefs are “not good” as evaluated with reference to epistemology’s standards: in the jargon of epistemology, they are *unjustified*, or *unwarranted*, or *irrational*, or *unreasonable*, or It is natural to suppose that this fact alone establishes that you are *not permitted* to believe as you do, that you are *required* to adjust your beliefs so that they avoid being “not good” from the epistemic point of view. Moreover, it seems to many that these claims of what is permitted and required are not contingent on what you happen to think about the value or desirability of epistemically good beliefs. If these appearances are to be trusted, then the evaluative standards of epistemology – unlike the evaluative standards of fashion, dance, effective criminality, and landscape architecture – are associated with categorical normative verdicts. In that case, we face a philosophical challenge: to vindicate the categorical nature³ of epistemically normative claims regarding what one is permitted or required to believe.

I would love to be able to offer an account of this sort. Unfortunately, I do not yet know how to do so. Instead, I will aim at something a bit less ambitious. My claim will be that on a great many occasions, epistemically normative claims regarding what one is required or permitted to believe can be generated on the basis of what we owe one another as social creatures. My argument will fall short of establishing that this is always true, for all beliefs at all times. But it will establish something that I hope will be of interest nonetheless: in those cases in which such an epistemically normative claim holds, the requirement or permission is not contingent on what one happens to want or what one happens to value. This result should be of some interest if only because it opens the prospect for a generalization of my argument: *if* it can be established that the sorts of consideration that I will be identifying here hold generally – something I will not attempt to show – then we will have the basis for a thoroughly social approach to the nature of epistemic normativity. In what follows, then, I will argue that there are cases in which there are epistemically normative claims regarding what one is required or permitted to believe, where this reflects what we owe one another as social creatures.

3. The Basic Argument

The sorts of case I will be discussing, then, are those in which (I maintain) we owe it to one another to adjust our doxastic attitudes in a manner that conforms to the evaluative standards of epistemology. To establish this claim, I will be focusing, in the first instance, not on the believing subject herself but on other parties, and in particular on what others are entitled to expect of the believing subject herself. The argument I will be developing is that there are a wide range of cases in which we are entitled

²Many people, though not everyone: there are consequentialists, whose view is that subjects ought to conform to the standards of epistemology because in so doing we maximize expected (epistemic) utility.

³Or something in the spirit of the categorical nature; see below.

to expect each other (*qua* believers) to conform to the evaluative standards of epistemology. My Basic Argument is as follows.

Basic Argument

Premise 1 For all subjects S and T, if T is entitled to a normative expectation that S will Φ , and there are no considerations that defeat or undermine this entitlement, then S ought to Φ .

Premise 2 Given epistemic subjects S and T, there are many cases in which (i) T is entitled to a normative expectation that S conform to evaluative epistemic standards in belief-formation and belief-management, and (ii) there are no considerations that defeat or undermine this entitlement.

Conclusion There are many cases in which epistemic subjects ought to conform to evaluative epistemic standards in belief-formation and belief-management.

Premise 1 of the **Basic Argument** should not be particularly controversial, though in the next section I will say a few words in its defense anyway. The real work will be in defending Premise 2. I will be presenting one main argument for it – what I call the ‘argument from cooperation.’ After that I will present (but not develop) one possible generalization of the argument I present.

4. Defending Premise 1 of the Basic Argument: entitled expectations of others

What can be said in defense of **Basic Argument**’s Premise 1, the claim that for all subjects S and T, if (i) T is entitled to a normative expectation that S will Φ , and (ii) there are no considerations that defeat or undermine this entitlement, then S ought to Φ ?

To address this let me begin by making clear what I have in mind with my terminology here. By speaking of a subject’s being *entitled* to expect such-and-such I am speaking of the subject’s being *permitted* to expect such-and-such, where the nature of the permission determines the sort of entitlement at issue. Given my aim in this paper, I will restrict myself to entitlements of an interpersonal nature – whether legal, social, moral, or interpersonal in some other sense. (The **Basic Argument** should be interpreted accordingly.) Thus a legal entitlement is a legal permission e.g. to do such-and-such: one who is legally entitled to do such-and-such has the legal right to do so, so that she could not be condemned on legal grounds for doing so. A social entitlement is a social permission e.g. to do such-and-such: one who is socially entitled to do such-and-such has the social right to do so, so that she could not be condemned on social grounds for doing so. And a moral entitlement is a moral permission e.g. to do such-and-such: one who is morally entitled to do such-and-such has the moral right to do so, so that she could not be condemned on moral grounds for doing so. Premise 1 talks of an entitlement to a *normative* expectation. Here I mean to be contrasting this sort of expectation with an empirical or predictive expectation, which in effect is an ordinary belief about the future course of experience (and which accordingly can be assessed with respect to the evaluative standards of epistemology). Normative expectations, by contrast, are attitudes whereby we hold others responsible. Thus one might normatively expect one’s children to do their chores (in the sense that one holds them responsible for doing so). Such an expectation can be appropriate even when one has ample evidence that they won’t do so.

I move next to the second condition in the antecedent of Premise 1, the claim that there are no considerations that defeat or undermine T’s entitlement to the normative

expectation that $S \Phi$. Here I have in mind the range of things that would excuse S from Φ ing: conflicting duties only one of which can be satisfied, other more pressing (legitimate) practical matters that prevent S from Φ ing, and so forth. For purposes here, I am treating these as defeating T 's entitlement to the normative expectation that $S \Phi$: if S is excused from Φ ing, then T isn't entitled to expect S to Φ . This does not mean that T isn't owed some sort of recompense in those cases where, having been entitled to expect that $S \Phi$, this entitlement is defeated by considerations that excuse S from Φ ing. Perhaps T is owed some sort of recompense in some cases. The point is simply that when S is excused from Φ ing, T is no longer entitled to expect S to Φ . The stipulation that there are no conditions defeating or undermining T 's entitlement to expect S to Φ , then, is meant to ensure that there are no excusing conditions present.

Given these understandings of the terms involved in Premise 1, this premise says in effect that when T has (legal or social or moral or ...) permission to hold S responsible for Φ ing and S has no excuses for not Φ ing, S ought to Φ . Although this should be obvious, I will nevertheless defend it. I will do so by considering three cases, according to whether the entitlement/permission in question is legal, moral, or social. My conclusion will be that in each case, Premise 1's consequent holds when the antecedent conditions hold, thereby vindicating the conditional itself.

Construed in legal terms, Premise 1 asserts that for all subjects S and T , if T is *legally* entitled to expect S to Φ , and there are no considerations that defeat or undermine T 's permission in this regard, then S ought to Φ . This claim seems obviously true. Suppose you have the legal permission to expect me to Φ , and nothing undermines or defeats your permission. Then you have the legal right to hold me responsible for Φ ing. But it would seem that the only thing that could give you that right is my having the corresponding legal duty to Φ . (How else could you have such a right?) And in that case, since your right to expect this of me was not defeated – and so was not defeated by e.g. a moral duty not to Φ – I ought to Φ . On a construal on which the entitlement in Premise 1 is a legal one, this premise is vindicated.

Consider next a construal on which the entitlement in question is a moral one. Premise 1 asserts that for all subjects S and T , if T is *morally* entitled to expect S to Φ , and there are no considerations that defeat or undermine T 's permission in this regard, then S ought to Φ . This claim would appear to be a trivial implication of the nature of legitimate moral expectations of others. Suppose you enjoy a moral entitlement to expect me to Φ . And suppose that there are no considerations that defeat or undermine your permission in this regard: this is no case of conflicting *prima facie* moral duties etc. Then you have the moral right to hold me responsible for Φ ing. But it would seem that the only thing that could give you that right is my having the corresponding moral duty to Φ . And in that case, I ought to Φ . The 'moral' version of Premise 1 is vindicated as well.

Finally, consider Premise 1 where the entitlement is construed as a social entitlement. To get a sense of what sort of entitlement this might be, consider the sort of entitlement you enjoy when you participate in a (legitimate) rule-governed social practice. Suppose that there is a widespread practice in your community of garbage pick-up on Wednesday, where the rules are that if you want your garbage picked up, you need to bring your (full) bins to the curb and have them face outward by no later than 6 a.m. Wednesday morning. If you do so, you are entitled to expect that your garbage will be picked up. The expectation here is normative: after all, you are not (merely) predicting that your garbage *will* be picked up if you have your bins ready by the noted time; you are *normatively* expecting pick-up under these circumstances. At the same time, you

entitlement to this normative expectation would not appear to be *legal* or *moral* in nature. What then is the nature of the entitlement? I call it a *social* entitlement since it seems to me that the entitlement itself derives from the nature of participation in a legitimate social practice itself: if you are a participant and the practice is a legitimate one, then you are entitled to expect that the other participants follow the rules of the practice. (In Goldberg (2017) I called these ‘practice-generated’ entitlements, but I think ‘social’ will do for now.)

Interpreted in terms of ‘social’ entitlements, Premise 1 of the **Basic Argument** says that for all subjects S and T, if T is *socially* entitled to expect S to Φ , and there are no considerations that defeat or undermine this entitlement, then S ought to Φ . So suppose that there is a legitimate social practice with a given rule Γ , that S and T are both participants in this practice (and this is a piece of mutual knowledge). Then T is socially entitled to expect S to follow Γ . And suppose too that there are no considerations that defeat or undermine this entitlement: that is, there are no emergencies or conflicting demands (and, in general, no overriding practical considerations) that would prevent S from following Γ (and in this way defeat one’s social entitlement to expect S to follow Γ). Under such conditions, S ought to follow Γ . Moreover, the point is general. Whenever there is a social entitlement to expect that S do such-and-such, and there are no considerations that defeat or undermine this entitlement, S ought to do such-and-such. This is ensured by the following: the social entitlement to expect derives from the practice itself, as any legitimate practice entitles its participants to hold one another to the rules themselves. And if this is so, then so long as there are no considerations that defeat or undermine this entitlement to hold another to the rules, if T is socially entitled to expect S to Φ , this is because S and T are participants in a legitimate social practice with rules Γ which mandate that S Φ , and so S ought to Φ . In short, the ‘social’ version of Premise 1 holds as well.

This concludes my defense of Premise 1 of the **Basic Argument**.

5. Defending Premise 2 of the Basic Argument: the Argument from Cooperation

Premise 2 of the **Basic Argument**, by contrast, is controversial. It says that given epistemic subjects S and T, there are many cases in which (i) T is entitled to expect S to conform to evaluative epistemic standards, and (ii) there are no considerations that defeat or undermine T’s entitlement. I want to offer one core argument for this claim, and then suggest one way it might be extended. The main argument I will call the argument from the nature of cooperation; the potential extension will appeal to what I will call “negative epistemic externalities.”⁴

The main argument I want to give for

Premise 2 Given epistemic subjects S and T, there are many cases in which (i) T is entitled to a normative expectation that S conform to evaluative epistemic standards in belief-formation and belief-management, and (ii) there are no considerations that defeat or undermine this entitlement.

is the argument from the nature of cooperation. The idea is that whenever we are entitled to expect cooperation from others, we are entitled to expect them to conform to evaluative epistemic standards. I might formalize the argument as follows:

⁴This argument was inspired by Cristina Bicchieri’s (2006) work on norms. She distinguishes social norms from regularities and customs, and she claims that one of the conditions under which real social norms emerge is when there are negative externalities involved in non-participation.

The Argument from Cooperation

CP₁ For all subjects S and T, whenever T is entitled to expect S to be cooperative, T is entitled to expect that any beliefs of S relevant to their cooperative efforts conform to the evaluative standards of epistemology.

CP₂ Given subjects S and T, there are many cases in which (i) T is entitled to expect S to be cooperative, and (ii) there are no considerations that defeat or undermine T's entitlement.

Therefore

C Given epistemic subjects S and T, there are many cases in which (i) T is entitled to expect S to conform to evaluative epistemic standards, and (ii) there are no considerations that defeat or undermine T's entitlement.

I will spend the majority of my time here arguing for CP₁, which I suspect will be the more controversial of the claims. I do so by appeal to a point that Karen Jones has made in connection with cases in which one is entitled to *trust* another person; my argument will be that what she says of trust holds, more generally, of cases in which we are entitled to expect cooperation (or at least that another be cooperative); and that in such cases our entitlement to expect cooperation involves an entitlement to expect conformity to the evaluative standards of epistemology (regarding any belief relevant to the cooperative effort).

Arguing that the norms of trust and trustworthiness are those that “fit us for participation in trust relationships” (Jones 2017: 102), Jones points out that the basis for others' expectations of us in this regard reflect “the value of enabling the extension of our agency by the non-coercive recruiting of the agency of others to remedy our finitude” (Jones 2017: 103). Her explanation of this point is worth quoting at length:

As finite social creatures other agents are a particular salient source of risk to us, but they also provide a remedy for our finitude, for together we can do what we cannot do alone, whether because the activity itself is necessarily a shared one (waltzing), or because it requires divisions of time, labour, and skill (most of the activities in our daily lives). Our reflectiveness provides a remedy, by enabling us to recruit our sociality to solve the problem of our finitude. Because we have a theory of mind, we can make decisions that take into account the mental life of others. There is thus available to us a distinctive way of responding to the fact of other agents' dependency *through recognizing that very dependency*. ... This dual structure of dependency – counting on the other, in a domain, and counting on them to respond to the fact that we are counting on them – is the heart of trust. (Jones 2017: 99–100; all italics in original)

She goes on to write that

Trust when met with trustworthiness allows us to enhance the effectiveness of our agency. As finite social agents we have a pressing interest in being able to do this. However, the interest that we have in being able to extend our agency in this way is not a distinctly moral interest. It is about making agency *effective* by drawing on the agency of others rather than about making it *good*. Extended agential power is an ends-independent value. Whatever our ends, we want to be able to recruit the

agency of others: even counter-ethical projects require recruits. (Jones 2017: 101–2; all italics in original)

I think Jones is on to something extremely important here, and it is worth unpacking it a bit.

I want to borrow from Jones' view a focus on the norm-generated expectations we have of others when we engage in trusting activities with them, and the nature of the expectations we have under these circumstances. I offer the following gloss and expansion⁵ of her view. We are deeply and ineliminably social creatures. This is seen in our actions (among other things). By oneself, one can only do so much. But with the cooperation of others, we can do much more.⁶ Of course, the move to cooperate with others opens us up to certain risks. For one thing, we can be taken advantage of when others have hidden motives. But even when others are sincerely trying to help in our cooperative efforts, we expose ourselves to the risk of failure if others' efforts are *incompetent*. Insofar as our engagements with others are of this trusting variety – for Jones this means that we are counting on them, and in addition we are counting on them to respond appropriately to our counting on them – we form normative expectations of them. When our trust is legitimate, we are entitled to these expectations. The expectations in question are of trustworthiness, where to be trustworthy is to be responsive to another's trust, and where this in turn involves not only sincerity but also competence – including doxastic competence.⁷ According to Jones this entitlement is not a moral one; it is rather a social one, deriving from the generic value of (cooperative) practical agency. Accordingly, as Jones notes, it is one to which we are entitled even in cases in which the ends towards which we are cooperating are morally dubious.⁸

Although Jones herself did not say this, I want to emphasize that the expectation of trustworthiness includes an expectation (not only of sincerity, but also) of competence – where this includes a doxastic dimension. To illustrate, suppose you expect my trustworthiness in connection with my Φ ing. And suppose further that, recognizing the legitimacy of your expectation of my trustworthiness, I sincerely aim to vindicate your expectation, and so aim to Φ . If I have false beliefs about when or how to Φ , I am likely to fail to Φ altogether or fail to Φ at the appropriate time; and even if I do Φ at the appropriate time, I am likely to do so poorly (in a way that does not vindicate your expectation of my trustworthiness).⁹ Consequently, if I am to Φ in a way that does vindicate your expectation, I must be *relevantly doxastically competent* – I must have whatever true, well-grounded beliefs that are (necessary and/or) sufficient for Φ ing in the relevant way.¹⁰ So your expectation of my trustworthiness implicitly involves an expectation of my relevant doxastic competence. It is here, of course,

⁵I will flag the expansions in footnotes.

⁶I speak of cooperation generally, whereas Jones speaks only of situations involving trust. See below.

⁷The claim that trustworthiness entails competence, including doxastic competence, is my expansion to Jones' own analysis. In defense of this expansion I note, first, that it is standard to regard trustworthiness as including both a sincerity dimension and a competence dimension, and second, that it should be obvious that one who is doxastically incompetent – whose beliefs are unjustified – will not be relevantly competent. See the following paragraph in the text.

⁸This doesn't mean that the ends towards which we are cooperating are beyond moral assessment!

⁹This might appear to assume that knowledge-how can be reduced to knowledge-that. But all that my argument requires is that some false beliefs regarding when or how to Φ can result in incompetence at Φ ing. Since this can hold even if the reductionist thesis is false, I am not assuming the reduction.

¹⁰Why *well-grounded* rather than merely true? Because competence requires as much: if the truth of one's beliefs regarding how or when to Φ is a matter of luck, then one can't be said to Φ competently (as opposed to luckily).

that we see the expectation to satisfy the evaluative standards of epistemology. If I fail to Φ (or fail to Φ in a way that vindicates your expectation of my trustworthiness), where this is the result of beliefs I have about when or how to Φ which are themselves based on inadequate evidence, then clearly you have a claim on me: in being entitled to expect my trustworthiness, you are entitled to expect me to avoid this sort of situation.

Jones' view is that this entitlement is not a moral one. This is important, as it ensures that our trust-related expectations of one another (in contexts in which trust and trustworthiness are at issue) hold independent of the ethical dimensions of our practical engagements with one another. As a result, others would have a claim on us *whenever* we violate the evaluative standards of epistemology on a topic relevant to the success of our cooperative engagements. This should be obvious when one's failure to satisfy those standards results (or would result) in the failure of the cooperative project itself. But the point holds more generally: so long as one fails to satisfy the evaluative standards of epistemology on a topic relevant to the success of our cooperative engagement, one's doing so *unduly risked* the failure of that project. The argument bears on any beliefs whose truth *potentially* bears on the success of the activity itself.

While I myself am inclined to agree with Jones that one can be entitled to expect trustworthiness in others even in contexts of illicit or immoral behavior, I recognize that many will disagree.¹¹ Happily, my argument does not need to endorse Jones' point in this regard. What my argument requires is that in contexts in which we are entitled to expect trustworthiness, we are entitled to expect doxastic competence. If Jones is correct in her contention that the entitlement to expect trustworthiness, based on the need to extend our agency, extends even to situations that involve immoral activity, then this will extend the reach of my argument to situations involving immoral activity. However, if she is wrong – if our entitlement to expect trustworthiness in others does not extend to situations in which we are engaging in illicit activities with them – then my argument will not extend to situations that involve immoral activity. In that case, the scope of my argument will have been modestly affected. But the argument itself would still go through (in all situations in which we continue to enjoy the entitlement to expect trustworthiness). I remain neutral on whether she *is* correct on the scope of the entitlement to expect trustworthiness.

I conclude, then, that if (bracketing the question of scope) Jones is correct in the rest of what she says about the norms of trust and trustworthiness, then we are entitled to expect doxastic competence, and so we are entitled to expect conformity to the evaluative standards of epistemology, whenever we are entitled to expect trustworthiness from another. (In that case the argument applies to any and all beliefs whose truth potentially bears on the success of the activity in which there are relations of trust.) But I think Jones' point holds more generally than this. In particular, I think it holds with respect to *any cooperative activity* whatsoever (whether it involves trust in Jones' sense or something less than that).¹² My core contention, then, is this:

COOP If we are entitled to expect cooperation, then we are entitled to expect conformity to the evaluative standards of epistemology by those with whom we are cooperating.

¹¹I thank David Henderson for indicating the need to address this point.

¹²For the significance of the difference between the general class of cooperative activity and the more restrictive class of trust-involving activity (where trust is understood in Jones' more restrictive sense), see Nickel (2017).

The basis on which Jones articulates the norms of trust and trustworthiness – namely, our general interest in extending the reach of our agency to include the efforts of others – holds not only in relations of trust but in relations of cooperation more generally.¹³

I now want to argue for CP₂, the claim that there are cases in which, given subjects S and T, T is entitled to expect S's cooperation, and there are no considerations that defeat or undermine T's entitlement. In what follows I will focus on the first part of this claim, asserting the existence of cases in which T is entitled to expect cooperation; having established this, I submit, is sufficient to establish CP₂ on the (imminently plausible) assumption that at least some of these cases involve no defeaters.

There are a variety of different ways one might come to be entitled to expect cooperation, or at least that others be cooperative. Most obvious are the cases in which cooperation is freely offered by another. But there are also a variety of circumstances in which the cooperation of others is legitimately presumed, where it is only if one "opts out" that this presumption is cancelled. In this category I would place cases in which (i) one is in a situation in which it is mutually manifest that one is a member of a team or collective at least some of whose other members are present; (ii) one is planning with others; (iii) one is involved in conversation (here we think of Grice's Cooperative Principle);¹⁴ (iv) one is involved in a situation involving significant coordination with others, where others behave in ways that are widely regarded as indicating that they are participating in the coordinated efforts; or (v) one is in a situation in which local cultural (or perhaps familial or friendship) norms have evolved in which cooperation is expected and presumed, and no conditions obtain to defeat the presumed expectation. No doubt there are many others. Indeed, if Jones is right, contexts of this sort constitute "most of the activities in our daily lives" (2017: 99).

It is important to appreciate that the foregoing argument bears not only in circumstances in which we are entitled to expect another's *cooperation*, but also in circumstances in which we are entitled to expect another's *cooperativity*. By the latter I have in mind an expectation that we might have *in advance* of any particular cooperative action we have already agreed to initiate with them or in which we are already engaged with them. The expectation in question is that they are *available* as "candidate[s] as cooperative agents" (Searle 1990) for a range of would-be ventures involving coordination or joint action; and this is an expectation to which we are entitled *prior to their giving any indication of their willingness to do so*. Consider: even if one speaker has never engaged another speaker in conversation before, the one can presume on the other's cooperativity in this regard (a point made by Grice); even if one driver has never before encountered another particular driver, if they are both driving their cars on the same two-lane road but in opposite directions, they can presume on each other's cooperativity in coordinating safe passage for each other; and so forth. If this is correct, then our argument above holds, not just for circumstances in which one is entitled to expect another's cooperation, but for circumstances in which one is entitled to expect another's *general cooperativity* – that is, in which one is entitled to expect them to be cooperative (prior to their giving any indication of their being so).

There is an important limitation to the argument on offer so far: it bears only on those beliefs whose truth or falsity is relevant to the success of the would-be cooperative endeavor. But the appeal to the expectation of cooperativity enables us to extend the class of beliefs regarding which we are entitled to expect one another to conform to the standards of epistemology. In particular, the beliefs regarding which others are entitled to expect one to conform to the standards of epistemology are not only

¹³See Lahno, "Trust and Collective Agency" (2017).

¹⁴For extensive discussion, see Goldberg (2020).

those beliefs whose truth is relevant, or potentially relevant, to the success of a cooperative activity in which others are already engaged with one, but also to any belief whose truth is potentially relevant to *any potential cooperation* with one on any matter on which *cooperativity itself is or would be properly expected*. I do not know whether this is every belief whatsoever; but I do know that it is a good many of our beliefs. Since our lives are replete with domains in which cooperativity is properly expected, and since in any situation in one of these domains we are entitled to expect relevant doxastic competence in all matters relevant to the sort of cooperativity at issue (and so are entitled to expect others to conform to the evaluative standards of epistemology), our lives are replete with situations in which we are entitled to expect others to conform to the evaluative standards of epistemology.

One other point is worth emphasizing here, concerning a constraint on the defeat condition. In any situation in which people are entitled to expect cooperation, no one can get out from under the relevant normative expectations merely because he doesn't want to be relied on. On the contrary, when cooperation is legitimately expected, the entitlement to this expectation is defeated only by conditions that a reasonable person would recognize as a reasonable basis for failure to comply.¹⁵ This point is important since it makes clear that, on the view on offer, categorical epistemically normative claims are not defeated by an epistemic subject's mere desire to opt out. The normative claims themselves can hold even when a subject would prefer that they were not under the relevant requirement.

In short, Premise 2 of the **Basic Argument** is supported by considerations from our entitlement to expect each other's cooperativity, and the **Argument from Cooperation** is complete.

6. Defending Premise 2 of the Basic Argument: the Argument from Epistemic Externalities

I just offered an argument on behalf of the second premise of the **Basic Argument**:

Premise 2 Given epistemic subjects S and T, there are many cases in which (i) T is entitled to a normative expectation that S conform to evaluative epistemic standards in belief-formation and belief-management, and (ii) there are no considerations that defeat or undermine this entitlement.

The argument was from our entitlement to expect cooperativity from others. In this section I want to suggest how this argument might be extended further by introducing a new tool in our social epistemology toolkit. The tool is the concept of a negative epistemic externality, and I will use this tool in presenting the outlines of an argument that tries to extend the **Argument from Cooperation**.

In economics, an externality is "a side effect or consequence of an industrial or commercial activity that affects other parties without this being reflected in the cost of the goods or services involved" (quoted from Google's online dictionary). Investopedia characterizes it as

a consequence of an economic activity experienced by unrelated third parties. ... Externalities occur in an economy when the production or consumption of a specific good impacts a third party that is not directly related to the production or

¹⁵Insofar as the "reasonable person" standard is evaluative, this begs no question in the context of my proposed account of this subclass of categorical epistemically normative claims.

consumption. (Cited on 8 December from <https://www.investopedia.com/terms/e/externality.asp>)

Some externalities are positive, e.g. the pollination of surrounding crops by bees kept for honey.¹⁶ Some are negative, e.g. the pollution from the industrial plant by the local river. I will speak of an “epistemic” externality as a consequence of a subject’s being in a certain doxastic state, experienced by someone other than the subject herself. More specifically, an epistemic externality obtains when the following conditions are met: (i) a subject S is in doxastic state-type Δ ; (ii) that S is in Δ has (practical and epistemic) consequences on subjects other than S; (iii) among those affected are some subjects who aren’t presently in Δ .¹⁷ This characterization of an epistemic externality does not speak of the “cost of the good or service,” and so does not represent such a cost as failing to capture the (dis)value of the consequences in question. Even so, the notion of an epistemic externality as characterized in terms of (i)–(iii) remains in keeping with the traditional notion of an externality: in the same way that a traditional externality is an indirect consequence (cost or benefit) of the production of a good or service which, though not reflected in the cost of the good or service produced, nevertheless affects third parties (who benefit from or bear the brunt of the consequence), so too an epistemic externality is an indirect consequence in the production or sustenance of a doxastic attitude by a given individual, which consequence affects many other people (who benefit from or bear the brunt of the consequence), including some who do not come to be in Δ (and so who in this sense don’t “pay the cost” of being in that state).

The claim that there *are* epistemic externalities – the claim that conditions (i)–(iii) are sometimes jointly satisfied – will be the basis of what I will call the **Argument From Epistemic Externalities**. In effect, this argument appeals to the existence of *negative* epistemic externalities to extend the argument above, claiming that each of us is socially entitled to expect that everyone conform to the evaluative standards of epistemology: this expectation embodies the attempt by social information-seeking creatures like us to manage the negative epistemic externalities of our systematic and ineliminable reliance on others. We can use this idea to present a variant on the argument from cooperativity; the point of the **Argument from Negative Externalities** is to illuminate what might underwrite our general entitlement to expect cooperativity. The argument here can be presented as a variant route to a strengthening of Premise 2 of the **Basic Argument**. This route can be represented as having two premises and a conclusion:

Argument from Epistemic Externalities

EE₁ If (a) Φ ing is something regarding which there are evaluative standards of assessment, (b) there are negative externalities involved in a subject S’s failing to meet those standards in Φ ing, and (c) those who suffer the costs of these negative externalities would have, in virtue of this, a legitimate (legal or moral or social) claim on S,¹⁸ then other subjects are entitled to expect that S conform to these evaluative standards.

¹⁶The example is from Google’s online dictionary.

¹⁷Obviously, ‘ Δ ’ designates a type of doxastic state – one that can be instantiated in more than one subject.

¹⁸I will speak of H as having a ‘moral claim’ on S when S failed to satisfy some expectation which H was morally entitled to have of S; and *mutatis mutandis* for H’s having a ‘social claim’ on S.

EE₂ For every subject S, conditions (a)–(c) hold when Φ ing = acquiring and sustaining doxastic attitude-type Δ , and the evaluative standards in question are those of epistemology.

Therefore

C For every subject S, other subjects are entitled to expect that S conform to the evaluative standards of epistemology.

The conclusion of this argument, C, is not quite Premise 2 of the **Basic Argument**: whereas Premise 2 itself is an existential claim, C is a universal claim; and whereas Premise 2 makes a claim about the existence of a range of cases in which there is an entitlement to a normative expectation under conditions in which there are no defeaters, C speaks only of the entitlement to a normative expectation (it is silent on the matter of defeaters). Still, if we can establish C, Premise 2 would follow on the (imminently plausible) assumption that some of the cases in which C holds are cases in which there are no defeaters present. In this way the **Argument from Epistemic Externalities** is an attempt to extend the conclusion of the **Argument from Cooperation**.

EE₁ should not be particularly controversial. The simple idea behind it is that others are entitled to expect you to meet the standard of goodness in a given domain when the following condition holds: the effect of your failure to meet that standard of goodness is that (i) others would be harmed and (ii) they would have a corresponding legal or moral or social claim on you.¹⁹ To illustrate, suppose that Driving Davis (henceforth ‘DD’) disregards traffic laws while driving, and so fails to meet some legal standard for driving a motor vehicle. And suppose further that in doing so DD causes damage to Victimized Velazquez’s (VV’s) vehicle. Merely in virtue of these facts VV would then seem to have a claim on DD. Of course, if the reason DD was disregarding the traffic laws was that VV was pursuing him with reckless abandon, flashing a gun in DD’s face, then that *prima facie* legal claim would be either illegitimate or defeated. But the fact that, absent DD’s having some legitimate excuse, VV would have a claim against him, suggests that VV is entitled to expect DD to have followed the proper driving standards (those embodied by the laws pertaining to motor vehicle operation).²⁰ EE₁ holds in such a case.

Of course, one example does not a general principle make. To establish the general principle, I must argue that EE₁’s consequent condition holds *whenever* its antecedent conditions hold. To this end I will need to be a bit more schematic in my presentation. So suppose that (a) Φ ing is something regarding which there are evaluative standards of assessment, (b) there are negative externalities involved in a subject S’s failing to meet those standards in Φ ing, and (c) those who suffer the costs of these negative externalities have, in virtue of this, a legitimate legal or moral or social claim on S. Given (b), some subject(s) will suffer from the negative externalities. Let T be such a subject. Given (c), T has legitimate legal or moral or social claims on S deriving from the harms she suffered from the negative externalities associated with S’s Φ ing badly. To have a legitimate legal or moral or social claim on S for Φ ing badly is to have a claim to the effect that S ought not to have harmed T in that fashion. (This claim holds so long as S has no

¹⁹(i) and (ii) here correspond to (b) and (c) from the first premise of the argument above, EE₁.

²⁰It is perhaps worth noting that the order of explanation may not be the one suggested by this argument; arguably, it is VV’s right to expect DD to have followed the standards that generates VV’s moral claim against DD. But this is no matter, as EE₁ does not purport to establish an explanatory claim; it is merely a conditional claim.

relevant excuse that defeats the claim.) Since the fashion in which T was harmed was itself the result of S's Φ ing badly, we can infer that S ought not to have Φ 'd badly. But then T is entitled to expect that, insofar as she Φ s, S ought to Φ properly or well – which, in effect, is the same as saying that she ought to conform to the evaluative standards governing Φ ing. And what goes for T goes for *any* subject who is a possible recipient of the potential harms associated with the negative externalities of S's Φ ing badly. This holds for any value of Φ itself. Since this just is EE_1 's consequent condition, EE_1 itself is established.

This brings me to EE_2 . This tells us that when it comes to S's having doxastic attitudes, conditions (a)–(c) of EE_1 hold. Here the relevant evaluative standards are those of epistemology. So the corresponding claims (a)–(c) would be these:

- (a*) S's having doxastic attitude-type Δ is something regarding which there are evaluative (epistemic) standards of assessment,
- (b*) there are negative externalities involved in a subject S's failing to meet those standards in having doxastic attitude-type Δ , and
- (c*) those who suffer the costs of these negative externalities would have, in virtue of this, a legitimate legal or moral or social claim on S.

Although I won't present the case for this at great length, arguably all three are true. (a*)'s truth is assumed for the sake of this paper, and in any case should not be particularly controversial: the standards of epistemology are evaluative standards for the assessment of a subject's doxastic attitudes. (b*) and (c*) will require more detailed defense.

Consider then (b*), the claim that there are negative externalities involved in a subject S's failing to meet those standards in having doxastic attitude-type Δ . It should be obvious that there are *some* cases in which this holds. To my mind, perhaps the best example of this is seen in the phenomenon of implicit bias (and racist or sexist belief more generally): when it shows up in one's doxastic attitudes, one's implicit bias against people of color, for example, can have all sorts of negative effects on those communities.²¹ And this is far from the only kind of example we can find. Consider a variant of the classic case of the ship owner described by Clifford. In this variant, the owner has evidence which calls into question the structural integrity of his ship, but, wanting to ensure that he continues to make money in his shipping business, unwarrantedly dismisses this evidence as misleading. So, continuing to believe his ship seaworthy, he sends it off to sea – where it sinks, and all people on it perish. Here, the ship owner fails to meet the evaluative standards of epistemology: he does not believe in accord with his evidence. And the result of his doing so are negative consequences that are borne by others (namely, the sailors and their families). These consequences were not intended, of course; they are indirect consequences of the owner's sending the ship out despite evidence calling the ship's seaworthiness into question. In this way, this case exemplifies (b*): there are negative externalities involved in failing to believe in accord with the evaluative standards of epistemology. It is also worth noting that it exemplifies (c*) as well: the families of the sailors have suffered a grievous harm, in virtue of which they have legitimate and undefeated social or moral claims on the ship owner.

Still, my claims in (b*) and (c*) are intended to be perfectly general: they purport to hold of *any* case in which a subject violates evaluative (epistemic) standards in having a

²¹This seems obvious when considering the (likely) behavioral consequences of such beliefs. Recently a stronger claim has been made: these beliefs harm the relevant communities independent of their (likely) behavioral consequences. See Basu (2019a, 2019b) and Basu and Schroeder (2019).

doxastic attitude. That is, I need to establish two further claims. The first is the Generalized Epistemic Externality Claim,

GEEC There are negative externalities involved *whenever* a subject S fails to meet epistemic standards in having doxastic attitude Δ .

The second is the claim, which I will label the ‘Subsidiary Claim,’ that as a result of GEEC the affected individuals have a (legal or moral or social) claim on S. To establish these one must do more than cite particular cases.

I aim to do so by citing two different considerations in defense of GEEC and the Subsidiary Claim. One of the considerations is practical: whenever S fails to meet epistemic standards in instantiating Δ , S thereby unduly risks putting others at risk²² of *practical* harms arising out of potential cooperative activity whose success potentially depends on true beliefs about the relevant matters. (It also imposes a cost on me, that of having to take the time and effort to confirm that those with whom I am engaging in cooperative projects are in fact relevantly doxastically competent.²³) The other consideration is epistemic: whenever S fails to meet epistemic standards in instantiating Δ , S thereby unduly risks putting others at risk of *epistemic* harms e.g. arising out of the testimony one would give on the matter in question. In both cases, I argue, the result is that those third parties affected by epistemic externalities have a (moral or social) claim on the subject S responsible for the externalities.²⁴ This is true even if *the risk* of these harms *does not materialize*. If this is so, then, just as communities have an interest in imposing regulations directed at managing traditional externalities, so too communities have an interest in imposing regulations directed at managing epistemic externalities. The regulations themselves involve the application of the evaluative standards of epistemology to the doxastic lives of those with whom we interact.

I begin first with the practical case for GEEC and the Subsidiary Claim. I submit that whenever a subject S fails to meet epistemic standards in having doxastic attitude-type Δ , S thereby unduly risks putting others at risk of practical harms arising out of any activity (joint or otherwise) which S participates in whose success depends on true beliefs about the relevant matters. I will call the latter – the risk of practical harms – the *first-order* risk, and I will call the former – the risk of exposing others to this first-order risk – the *higher-order* risk.

The need to speak of higher-order risk can be made clear by illustration. Suppose that the City of Chicago’s Department of Transportation (DOT) decides to start digging a huge hole in one of the lanes of Lake Shore Drive. Suppose that DOT does nothing to indicate to motorists that they are doing so. Then those motorists who drive along the relevant part of Lake Shore Drive are exposed to the (first-order) risk of serious injury. But suppose that it just happens (miracle of miracles!) that no one drives along Lake Shore Drive for the period during which DOT was working on the hole. (They remain ignorant of the risk, but luckily everyone decides to avoid the Drive.) Even so, insofar as the DOT had no adequate grounds for thinking that motorists would avoid the Drive,

²²The reason for my speaking of a higher-order risk – the risk of putting others at risk – will emerge below.

²³Economists have spoken of the cost of doing business in low-trust communities; this as I see it is a species of the same sort as I am characterizing here, which concerns the (cognitive and resource) costs of cooperation under conditions in which one is not confident of others’ relevant doxastic competence.

²⁴In what follows I will be moving back and forth between claims regarding others’ expectations that one not violate epistemic standards, and claims regarding others having a (moral or social) claim on one if one does violate such standards.

the DOT risked placing others at risk: it was guilty of doing something that generated this *higher-order* risk. This makes clear that even when few if any are exposed to the first-order risk, there remains the higher-order risk. That this is an externality is clear: insofar as people are entitled not to be exposed to the first-order risk (and when they are they have claims against those who would so expose them), it would seem that they are entitled not to be put in a situation in which only brute contingencies protect them against being exposed to that first-order risk. You are not let off the hook for not having built a fence around your swimming pool by the fact that, as a matter of fact, no one walked near your yard (and so no one was at any risk of drowning).

For their part, the first-order risks involved in the cases I am imagining – the risks that manifest the epistemic externalities I am highlighting – are easy to appreciate. I am assuming here that the standards of epistemology are standards of goodness in our efforts at acquiring truth and avoiding error, so that to fail to satisfy them increases the (objective or subjective) risk of falsehood.²⁵ Consequently, if you have beliefs that fail to satisfy epistemic standards – that fail to be justified – you are running an increased (first-order) risk of falsehood. The problem with doing so is that the success of a (solo or joint) project or coordinated effort often depends on the truth of relevant beliefs. In that case, if your unjustified beliefs are false you have damaged the prospects for a successful effort. But whether or not your unjustified beliefs are false, you have exposed others to an *increased (first-order) risk of falsehood*, and so an increased (first-order) risk of damaging the prospects for success in the (solo or joint) action or coordinated effort. Insofar as others can be affected by S's actions and the outcome of S's efforts even when they themselves are not participating in those actions with S, they have a claim on him. This claim is akin to the claim made above in connection with the **Argument from Cooperation**, extended here to cases in which one can be affected by the actions of others.

One particular externality is worth dwelling on at length: the testimonial spread of false or otherwise unwarranted belief. Were a subject S to violate the evaluative standards of epistemology, S would unduly put others at (first-order) risk of *epistemic* harms arising out of any testimony S would give on the matter in question. An initial point on this score is easy to appreciate. If your belief that p is unjustified and, nevertheless, you testify that p, then insofar as others' take your word for it, you have increased the (first-order) risk that they suffer from the epistemic harm associated with *acquiring a false belief*.²⁶ Once again, I submit that others are entitled to expect that you not do this.²⁷

²⁵'objective or subjective': this will depend on what the correct evaluative standards are. I am trying to remain neutral on this. In what follows I will refrain from adding the 'objective or subjective' modifier, though it should be understood to be in place when I speak of the increased risks of falsehood when epistemic standards are not met.

²⁶Here I note that anyone who cooperates with those who have accepted your false testimony is exposed to the sort of risks described above.

²⁷David Henderson (personal communication) suggests a dilemma, based on whether or not we assume that

(FB) The falsity of a belief is an unconditionally bad thing (from the epistemic point of view).

If we embrace FB, then we can appeal to the unconditional badness of false belief in order to ground epistemically normative claims, without having to appeal to what we owe one another. But if we reject FB, then it is hard to see how we can have an entitlement to expect that others not expose us to false belief in the first place. In response, I embrace the first horn, and suggest that the appearance of a problem reflects the conflation of the evaluative with the normative. Even if false belief is unconditionally bad from the epistemic point of view, what is wanted is a defense of the claim that we ought to avoid this (epistemically

The claim that others are entitled to expect a speaker S to avoid *attesting to* propositions which S herself is not justified in believing is a familiar one. It reflects the norm governing acts of testifying. For many theorists, these acts are acts of assertion, and the relevant standard is given by the norm of assertion.²⁸ For other theorists, there are distinctive norms governing the act of testifying (some prefer to speak instead of the act of telling).²⁹ At any rate, all appear agreed that the norm itself requires that the speaker S possess the relevant epistemic authority to testify. On most views this amounts to having the knowledge in question (and manifesting it in the testimony). Insofar as this is a speech act norm, others who observe an act of this type are entitled to expect S to have satisfied the norm in performing the act. And since it enjoins having the relevant epistemic authority, this means that others are entitled to expect S to have the relevant authority. Finally, since this requires having justification, others are entitled to expect S to be justified in believing the propositions to which they attest.

It is important to bear in mind that, for my purposes, the expectation to which audiences are entitled on this score is a normative one: we are entitled to hold the speaker responsible for being justified in the relevant sense. It is a further question what justifies the audience in *assuming or believing* that the speaker does in fact have the relevant justification.³⁰ No matter; whatever the answer to that question, it remains true that the audience who observes a speaker's testimony is entitled to the normative expectation that the speaker have the relevant justification required by the norm governing acts of that type.

Even with this qualification in place, however, one might wonder whether this argument is properly focused. For one might think that, for all this argument says, a speaker S can believe as she likes, so long as she doesn't *attest to* any proposition she fails to be justified in believing. What is wanted is an argument connecting the norms governing the speech act of testifying or telling to the expectations we have in connection with others' beliefs.

Some have argued for this connection by arguing that if belief is to play the various roles assigned to it, then the evaluative (epistemic) standards governing belief are – *must be* – the same as the standards articulated by the norm of assertion or testimony.³¹ Since I am dubious of the success of this sort of argument,³² I would like some other way of connecting the norm of testimony or assertion with the requirement that subjects satisfy the evaluative epistemic standards of belief. I submit that at a minimum there is an indirect relation between the speech act norm and the evaluative epistemic standards on belief. If one is to be competent at satisfying the norm governing acts of testifying, one must generally be able to tell when one satisfies the norm (and when not). I submit that if one were generally *unjustified* in one's beliefs regarding whether one satisfied the norm for proper testifying or asserting, then one would jeopardize one's competence in satisfying the speech act norm. So it would seem that if others are entitled to expect one to conform to the norm governing testifying, then at a minimum others are entitled to expect one to conform to the evaluative standards of epistemology for the class of beliefs one forms about *whether one satisfies the norm governing telling*.

unconditionally bad) outcome. Consider that etiquette might have some things that it considers unconditionally bad from the point of view of etiquette – that is, bad (etiquette-wise) in a way that does not depend on any (cultural or individual) contingencies – and yet we can still ask whether one ought to avoid those things.

²⁸See e.g. Fricker (1987), Sosa (1994), and Goldberg (2015).

²⁹See e.g. Fricker (2004), Moran (2006), Hinchman (2006), and Zagzebski (2012).

³⁰Contrary to Angus Ross (1986); see Goldberg (2015, 2020) for a response.

³¹See e.g. Williamson (2000: 255–6), Bird (2007: 95), Sutton (2007: 46), Stanley (2008: 52–3).

Another point is relevant here. Subjects cannot tell in advance which information will be worth attesting to, since they cannot tell in advance what informational needs their fellows will have. If this is so, then it will not do to adopt a policy of *believing as one will but ensuring that one's testimonies are proper*. For one will not know in advance whether one's belief will be on a subject on which others have informational needs one hopes to address. So it would seem that insofar as one aims to ensure that one's testimonies are proper, one ought to restrict one's beliefs to those that satisfy the evaluative standards of epistemology. And it would seem that others are entitled to expect as much, at least insofar as they are entitled to expect others to conform to the norm governing the act of testifying.

It is perhaps worth seeing this point as the analog, for testimony, of the practical point I borrowed from Jones above. Insofar as we are finite creatures, we rely on others to extend the reach of our agency. When others are entitled to expect this of us, they are entitled to expect us to satisfy the conditions on the sincerity and competence of our efforts – including our doxastic competence. The point above was that this holds true in any situation in which we are engaged in cooperative behaviors with others. But sharing information is one way to be so engaged. So this point is a special case of the more general point made previously. Even so, it is important to bring this point out separately, since the range of information which others might want and which we might be able to provide extends beyond the information that might be relevant to our (non-epistemic) practical cooperative projects. Once again, this supports both the Generalized Epistemic Externality Claim, GEEC, and the Subsidiary Claim.

This concludes the **Argument from Epistemic Externalities**. The idea behind this argument can now be appreciated: when one's beliefs fail to satisfy the evaluative epistemic standards of epistemology, one risks exposing others to the effect of one's own failed efforts (whether in solo or joint actions) whenever the efforts fail as a result of one's false or unwarranted belief. In effect, this is an "epistemic externality": it is an indirect consequence of a subject S's having an unjustified belief, where the effects of S's having such a belief bear on those beyond S herself. These effects are seen both in the cooperative practical projects in which S engages with others, but also in situations in which S is acting alone but where her efforts have effects on others. A special case of this sort of negative externality arises in the information-sharing practices in which S engages with others. In these cases, the negative effects of S's having an unjustified belief spill out to harm others, and result in their having (legal or moral or social) claims on S. But even when (in point of fact) others are not exposed to the (first-order) risks in question, those who exhibit poor epistemic conduct in belief-maintenance run the (second-order) risk of so exposing people. Just as we are entitled to expect that pool owners take precautions against accidental drowning by building fences around their pools *whether or not* others happen to go near their pool, so too we are entitled to expect that other subjects take precautions against exposing us to the risk of epistemic harms *whether or not* others happen to rely on them in relevant ways. The conclusion of this argument, then, is that each of us is entitled to normative expectations regarding other epistemic subjects in this regard, and so each of us is entitled to expect others to conform to the evaluative standards of epistemology – whenever another subject's actions expose others to this sort of (second-order) risk. Since we are social creatures by nature, most of our lives involve various forms of acting with others, and of actions that affect them; and so the scope of this expectation is wide indeed.³³

³³Below I consider a worry that arises in cases *not* involving cooperation.

7. Conclusion

In this paper I have argued that at least in some cases, and arguably in a great range of cases, epistemically normative verdicts – epistemology-driven verdicts regarding what subjects are categorically permitted or required to believe – can be seen to reflect what we owe one another as social creatures whose practical and epistemic dependence on each other is both systematic and ineliminable. To this end I tried to argue, first, that in contexts in which we are entitled to expect cooperation we are thereby entitled to expect our fellow cooperators to live up to the evaluative standards of epistemology, and second, that there is an epistemic analogue of a phenomenon familiar in the literature in political philosophy and economics (externalities). In effect, my claim that we are entitled to expect one another to conform to the evaluative standards of epistemology reflects our entitlement to expect others to be cooperative, and to expect them to manage their epistemic externalities.

Do these arguments – the argument from cooperation and the argument from epistemic externalities – show that (in the cases in question) the demand to conform to epistemic standards is *categorical*? Admittedly, this is not entirely clear. It may be that at best they show that *if* we are to be cooperative etc., then we are required to live up to those standards. Still, the argument would be of interest nevertheless. For one thing, some will find it plausible that there are cases in which there is a categorical requirement on us to be cooperative; and for such people, the argument will deliver the result that in these cases the requirement to live up to the standards is also categorical. But a weaker position is available: even if the demand to be cooperative isn't categorical in any case, it is plausible to think that this demand does not rest on what one happens to want or value. And if that much is true, then the arguments above will have succeeded in showing that, at least in the cases in which cooperation etc. is properly expected, the demand to live up to epistemic standards does not rest on what one happens to want or value. Finally, for Kantians I might offer one last speculation. Perhaps we should think of our disposition to cooperate as reflecting our social nature, part of who we are as a species – in which case grounding epistemic normativity in cooperation is no more to render it hypothetical than grounding moral normativity in our rational nature rendered it hypothetical for Kant.³⁴

References

- Basu R.** (2019a). 'What We Epistemically Owe to Each Other.' *Philosophical Studies* 176(4), 915–31.
- Basu R.** (2019b). 'The Wrongs of Racist Beliefs.' *Philosophical Studies* 176(9), 2497–515.
- Basu R. and Schroeder M.** (2019). 'Doxastic Wronging.' In B. Kim and M. McGrath (eds), *Pragmatic Encroachment in Epistemology*, pp. 181–205. New York, NY: Routledge.
- Bicchieri C.** (2006). *The Grammar of Society: The Nature and Dynamics of Social Norms*. Cambridge: Cambridge University Press.
- Bird A.** (2007). 'Justified Judging.' *Philosophy and Phenomenological Research* 74(1), 81–110.
- Fricke E.** (1987). 'The Epistemology of Testimony.' *Proceedings of the Aristotelian Society* 61, 57–83.
- Fricke E.** (2004). 'Testimony: Knowing Through Being Told.' In M. Sintonen, J. Wolenski and I. Niiniluoto (eds), *Handbook of Epistemology*, pp. 109–30. Dordrecht: Kluwer Academic Publishers.

³⁴With thanks to Amy Flowerree, John Greco, and David Henderson, for extensive written comments on earlier drafts of this paper; to Cristina Bicchieri, Grace Boey, Sinan Dogramaci, Peter Graham, Hannah Gunn, Klemens Kappel, Kareem Khalifa, Max Lewis, Coleen Macnamara, Lisa Miracchi, Carry Osborne, Carlos Pereira Di Salvo, Eric Schwitzgebel, Dan Singer, Mona Simion, Bob Talisse, Alessandra Tanesini, Howard Wettstein, and Kevin Zollman, for comments given in conversation; and to audiences at the University of Pennsylvania, the University of Nebraska, the University of Cardiff, the University of California-Riverside, and the University of Copenhagen, where I have given versions of this paper as a talk.

- Hinchman T.** (2006). 'Telling as Inviting to Trust.' *Philosophy and Phenomenological Research* 70(3), 562–87.
- Goldberg S.** (2015). *Assertion: On the Philosophical Significance of Assertoric Speech*. Oxford: Oxford University Press.
- Goldberg S.** (2017). 'Should Have Known.' *Synthese* 194(8), 2863–94.
- Goldberg S.** (2020). *Conversational Pressure: Normativity in Speech Exchanges*. Oxford: Oxford University Press.
- John K.** (2017). 'But I Was Counting on You!' In P. Faulkner and T. Simpson (eds), *The Philosophy of Trust*, pp. 90–109. Oxford: Oxford University Press.
- Lahno B.** (2017). 'Trust and Collective Agency.' In P. Faulkner and T. Simpson (eds), *The Philosophy of Trust*, pp. 129–48. Oxford: Oxford University Press.
- Moran R.** (2006). 'Getting Told and Being Believed.' In J. Lackey and E. Sosa (eds), *The Epistemology of Testimony*, pp. 272–306. Oxford: Oxford University Press.
- Nickel P.J.** (2017). 'Being Pragmatic about Trust.' In P. Faulkner and T. Simpson (eds), *The Philosophy of Trust*, pp. 195–213. Oxford: Oxford University Press.
- Ross A.** (1986). 'Why Do We Believe What We Are Told?' *Ratio* 1, 69–88.
- Searle J.** (1990). 'Collective Intentions and Actions.' In P. Cohen, J. Morgan and M. Pollack (eds), *Intentions in Communication*, pp. 401–15. Cambridge, MA: MIT Press.
- Sosa E.** (1994). 'Testimony and Coherence.' In B.K. Matilal and A. Chakrabarti (eds), *Knowing From Words*, pp. 59–67. Amsterdam: Kluwer Academic Publishers.
- Stanley J.** (2008). 'Knowledge and Certainty.' *Philosophical Issues* 18, 34–57.
- Sutton J.** (2007). *Without Justification*. Cambridge, MA: MIT Press.
- Williamson T.** (2000). *Knowledge and Its Limits*. Oxford: Oxford University Press.
- Zagzebski L.** (2012). *Epistemic Authority: A Theory of Trust, Authority, and Autonomy in Belief*. Oxford: Oxford University Press.

Sanford C. Goldberg is Professor of Philosophy at Northwestern University. He works on topics at the intersection of epistemology and philosophy of language. He is author of *Conversational Pressure* (Oxford University Press, 2020), *To the Best of Our Knowledge* (Oxford University Press, 2018), *Assertion* (Oxford University Press, 2015), *Relying on Others* (Oxford University Press, 2010), and *Anti-Individualism* (Cambridge University Press, 2007).