

The Relational Foundations of Epistemic Normativity

Cameron Boulton

Abstract: Why comply with epistemic norms? In this paper, I argue that complying with epistemic norms, engaging in epistemically responsible conduct, and being epistemically trustworthy are constitutive elements of maintaining *good epistemic relations* with oneself and others. Good epistemic relations are in turn both instrumentally and finally valuable: they enable the kind of coordination and knowledge acquisition underpinning much of what we tend to associate with a flourishing human life; and just as good interpersonal relations with others can be good for their own sake, standing in good epistemic relations is good for its own sake. On my account, we have reason to comply with epistemic norms because it is a way of respecting the final value of something that also tends to be an instrumentally valuable thing: good epistemic relations. Situating the account within the recent social turn in debates about epistemic instrumentalism, I argue that the dual-value aspect of good epistemic relations can explain important anti-instrumentalist intuitions, in a well-motivated way, within a broadly instrumentalist framework.

1. Introduction

The epistemic domain is comprised of distinctive *norms, reasons, and values*. One of the most basic questions we can ask about the epistemic domain is: why comply with epistemic norms, respect our epistemic reasons, and promote epistemic values? What explains the grip that epistemic norms, reasons, and values seem to have on us? Call this the *source question* about epistemic normativity.¹

My aim in this paper is to propose a novel answer to the source question. In my view, the answer lies in the quality of our relations with one another as epistemic agents. When agents have reciprocal sets of intentions, expectations, and attitudes that are oriented around the cultivation and utilization of their epistemic agency, they stand in a distinctively “epistemic relation”. Complying with epistemic norms, engaging in epistemically responsible conduct, and being epistemically trustworthy are all constitutive elements of maintaining good epistemic relations. Good epistemic relations are valuable in at least two ways. On one hand, they are instrumentally valuable: they enable the kind of coordination and knowledge-acquisition that underpins much, if not all, of what we tend to associate with a

¹ See the end of this introduction for further clarification on what I take to be at issue in the source question.

flourishing human life (on both an individual and collective level). But they also have *final* value. Just as standing in a variety of good interpersonal relations with others can be good for its own sake, standing in good epistemic relations is good for its own sake.² I will argue that these complex connections between epistemic norm compliance and the value of good epistemic relations (both instrumental and final) can be developed into a compelling explanation of our reason to comply with epistemic norms. Call this the “good relations view”.

One especially attractive feature of the good relations view is that it offers a satisfyingly *ecumenical* approach to the source question. As is well known, the debate between “instrumentalist” and “anti-instrumentalist” answers to the source question is longstanding. Powerful intuitions, such as those at play in the “too few reasons” problem (Côté-Bouchard 2015), are deeply recalcitrant and remain a significant site of disagreement. I will argue that the good relations view can accommodate key insights from both sides of this debate. While good epistemic relations are extremely useful, and our compliance with epistemic norms is important for that reason, there is a dimension of epistemic normativity that is best explained by the fact that good epistemic relations – like many other interpersonal relations, such as friendship – are finally valuable, and – as with other interpersonal relations – an important dimension of properly valuing them is by *respecting* rather than *promoting* their existence. In a slogan, we have reason to comply with epistemic norms because it is a way of respecting the final value of something that also tends to be an instrumentally valuable thing: good epistemic relations. I will argue that this dual aspect of good epistemic relations can account for important anti-instrumentalist intuitions, in a well-motivated way, within a broadly instrumentalist framework.

Here is the plan. In Section 2, I situate the good relations view within the source debate, motivating a distinction between “direct” and “social” forms of instrumentalism, and identifying my framework as being closely related to – but

² Following Christine Korsgaard (1983), I do not equate final value with “intrinsic value” – the value something has in virtue of its *intrinsic properties* – but rather the value something has as an end or for its own sake. Traditionally, “instrumental value” has been contrasted with “intrinsic value”. I agree with Korsgaard that this is misleading (Korsgaard 1983). In my view, the proper contrast of instrumental value is final value, and the proper contrast of intrinsic value is “extrinsic value”, or the value that something has in virtue of its relational properties.

importantly different from – the recent *social* turn in this debate. In Sections 3, 4, and 5, I offer an account of good epistemic relations, paying close attention to the features that make them both distinctively epistemic and finally valuable, and defending an account of how they are properly valued. In section 6, I argue that the good relations view has an attractive edge in dealing with two importantly different kinds of concern about social instrumentalist approaches – I focus on worries about the grip of epistemic norms in the absence of social frameworks, and empirically based considerations about epistemically flourishing social frameworks in the absence of epistemic norm compliance (Levy and Alfano 2020).³ In Section 7, I respond to a potential objection that the good relations view is uninformative or viciously circular. By pressing on a distinction between skeptical and explanatory approaches to the source question (which I draw immediately below), and between viciously circular and mutually supporting theoretical suppositions, I argue that it is not. In Section 8, I elaborate on the key motivation behind the good relations view highlighted at the outset. I argue that the structural and explanatory connection between the account’s claims about the instrumental and final value of good epistemic relations has important ecumenical advantages in comparison with a range of anti-instrumentalist approaches, specifically in its ability to address Christopher Cowie’s (2014) “argument from coincidence”. I also argue that the account enjoys independent motivation, focusing on the independent plausibility of its claims about final value. I close with a brief note on the role good epistemic relations already play in other areas of epistemology, such as the epistemic *harms* and epistemic *blame* literature. Putting all this together, a compelling case emerges for the relational foundations of epistemic normativity.

Before getting started, a couple of preliminary remarks about the source question are in order. There are different ways of framing the source question. We can ask it from a *skeptical* stance: that is, from the point of view of someone unconvinced about the grip of epistemic norms. We can also ask it from an *explanatory* stance: that is, from the point of view of someone who feels the grip of

³ The view can also help with a wider range of challenges. I focus on these two for sake of space, and because doing so showcases the structure of strategies that can be applied elsewhere.

epistemic norms, but wants to know what makes this the case. My interest is the latter stance. Sometimes it's put in terms of what "underwrites" our epistemic reasons (Worsnip 2022), or what "grounds" epistemic normativity (Côté-Bouchard 2016). Kurt Sylvan (2020) calls theories attempting to answer this question "foundational first-order theories of epistemic normativity" (2020: 3, 27). These are all ways of framing the source question from an explanatory stance. In what follows, I engage with the source question in terms of explaining the grip or force of epistemic norms. By "epistemic norms" I mean norms such as those enjoining us to proportion our beliefs to the evidence, assert only what we know, or be mindful of our biases. Rather than taking a substantive position on the correct epistemic norms, I use "epistemic norms" as shorthand for any normative dimension of the epistemic domain that seems to have the kind of grip or force epistemologists are interested in. By "grip" or "force", I have in mind what James Willoughby (2022) calls "a robust normative pressure that permeates our lives" (8). This need not entail that we have reason, always and everywhere, to comply with epistemic norms. But I take it that our reason to comply with epistemic norms is highly stable and robust. These points can surely be further precisified, but a more precise articulation is not necessary for present purposes. What matters is being able to explain, to the extent that there is reason to comply with epistemic norms, why this is so.

2. Instrumentalism: Direct, Anti-, and Social

The most sophisticated and worked-out answers to the source question tend to divide along instrumentalist and anti-instrumentalist lines. Early on, instrumentalism took a "direct" form, according to which, for any agent *S*, epistemic normativity gets a grip on *S* in virtue of standing in a certain relation to *S*'s goals, aims, or interests. We can further divide views on this relation into three categories: token epistemic-goal views,⁴ general epistemic-goal views⁵, and direct practical goal views.^{6 7} One of the core motivations of direct instrumentalism is that its success

⁴ Buckley 2021; Kelly 2003.

⁵ Foley 1987.

⁶ Kornblith 1993, 2002; cf. Carter 2022; Steglich-Petersen 2018.

⁷ See Worsnip 2022 for similar categorization.

would subsume epistemic normativity under a much broader class of normativity – instrumental normativity – and, according to defenders, thereby unify our meta-normative theorizing in powerful ways.⁸ The core difficulty for all forms of direct instrumentalism is the so-called “too few reasons” problem (Côté-Bouchard 2015). There seem to be clear cases of agents who have reason to comply with epistemic norms, even though doing so will not promote any kind of goal they plausibly have, or may even frustrate their goals. Given direct instrumentalism’s explanation of the source of epistemic normativity in terms of a connection to our goals, it seems unable to make sense of such cases.

The too few reasons problem motivates many to reject instrumentalism outright, embracing “anti-instrumentalism” in its place. This can take different forms – some argue that the source of epistemic normativity is constitutively related to the nature of belief,⁹ the etiological function of belief,¹⁰ and some think it is simply an explanatorily basic phenomenon.¹¹ I will not get into the details of these views (and corresponding challenges), since my aim is the positive one of motivating an alternative approach to abandoning instrumentalism altogether. I will, however, make some claims about the comparative explanatory power of my preferred approach over existing forms of anti-instrumentalism in Section 8.

An alternative approach to abandoning instrumentalism altogether is to reframe the connection between epistemic norm compliance and the promotion of our goals or aims. Rather than framing the connection as a *direct* one between epistemic norm compliance and an individual’s goals or aims, we can also frame the connection as being mediated by *social* phenomena. Rather than being a direct way of promoting one’s goals, perhaps compliance with epistemic norms is just a constitutive element of certain social structures or practices which are themselves a precondition on, or integral to, individuals’ and their communities’ flourishing.¹²

⁸ See Cowie 2014 and Worsnip 2022 for critical discussions of this and other ways of motivating direct instrumentalism.

⁹ McHugh 2012; Shah 2003; Shah and Velleman 2005;

¹⁰ Bird 2007; Burge 2003; Simion 2019, cf. Côté-Bouchard 2016.

¹¹ Kelly 2003; Wedgwood 2007; Parfit 2011; Worsnip 2022; Cf. Cowie 2014 and Sharadin 2018.

¹² Dogramaci (2012), Dyke (2021), Grimm (2009), Graham (2015), Goldberg (2018; 2020), Henderson (2020), Wei (forthcoming).

Two especially clear and compelling examples of social instrumentalism have recently been developed by Matthew Chrisman (2022) and Michael Hannon and Elise Woodard (forthcoming). For Chrisman, support for social instrumentalism lies in the role that compliance with epistemic norms plays in maintaining our “epistemic reputations”. He suggests that maintaining one’s epistemic reputation:

“[...] arguably, is itself crucial for securing epistemic support via testimony from other people. The idea, then, is this: since success in satisfying one’s own desires and interests so often depends on having such support from other people, individuals have a strong and ongoing reason to maintain their reputation as reasonable and reliable believers. This is why they have a strong default reason to care about whether their beliefs comply with epistemic norms (2022, 128).¹³

According to Hannon and Woodard:

...[I]t is not the particular interests or goals of any believer that give rise to an obligation to comply with epistemic norms. Our epistemic obligations arise from a less personal, more intersubjective perspective. More specifically, we are bound by the norms and expectations that structure a practice of mutual epistemic accountability, which promotes epistemic rule-following across the community. To flourish in a society, people need to cooperate in joint activities and coordinate expectations; in order to achieve these ends, we need to be able to expect others to keep agreements and comply with social conventions around which we can coordinate behavior. [...] This is why epistemic norms “apply to” or “bind” us all. (forthcoming, 12).

There are important differences between these claims. But unifying them is the idea that compliance with epistemic norms facilitates the kind of *knowledge acquisition* and *co-ordination* within epistemic communities that is integral to flourishing human life. This is the sense in which they’re forms of social instrumentalism. What is intriguing

¹³ Importantly, Chrisman also develops a “Rousseauian” view of the source of epistemic normativity, which he claims is ultimately more satisfyingly social. It may be that my own good relations view is closer to this proposal. I leave that question to the side in this paper.

about social instrumentalism is its promise in preserving the basic spirit of instrumentalism while avoiding the too few reasons problem. Even if there are cases where no individual goal would be satisfied by epistemic norm compliance, individuals should comply with epistemic norms because doing so is a necessary element of participating in a social framework that is integral to human flourishing (Hannon and Woodard, forthcoming).¹⁴

I am sympathetic to social instrumentalism, in large part because the too few reasons problem gives us strong reason to move away from direct instrumentalism; and I am not convinced by predominant anti-instrumentalist alternatives either (largely for reasons I will discuss in Section 8). I do, however, think social instrumentalism generates a host of further worries that have not been adequately addressed in the literature. The possibility of socially isolated agents, epistemic-free riders, and empirically grounded examples of communities who appear to flourish in virtue of *not* complying with epistemic norms (Levy and Alfano 2020), all raise important challenges for the framework. In my view, the key to adequately addressing these issues, and arriving at a more satisfying approach to the source question altogether, is a better understanding of *good epistemic relations*. In addition to facilitating the kind of knowledge acquisition and co-ordination that individuals and communities often need to flourish, participating in good epistemic relations is something worth doing *for its own sake*. The final value of good epistemic relations opens powerful conceptual space in the source debate, allowing us to deal with a range of challenges for social instrumentalism, and to accommodate anti-instrumentalist intuitions with materials that are structurally and explanatorily connected to a broadly instrumentalist framework.

3. Our Epistemic Relations with Others

What characterizes our epistemic relations with others? We can fill this out by thinking of epistemic relations as reciprocal sets of *intentions, expectations, and*

¹⁴ See Scott (2023) for skepticism about social instrumentalism's promise on this front. Since I'm myself not convinced existing social instrumentalist views fully satisfactorily answers the source question, I won't get into his criticisms here.

attitudes that are oriented around the cultivation and utilization of our epistemic agency (Boult 2021; forthcoming). Epistemic relations are related to ordinary interpersonal relations in that we can also helpfully think of many ordinary interpersonal relations, such as personal relationships, in terms of reciprocal sets of intentions, expectations, and attitudes. Consider friendship. In a sense, what it *is* to be friends with someone is to have certain reciprocal intentions, expectations, and attitudes towards that person and one's activities with them. At least in certain kinds of friendships, friends intend to be there for one another, to keep in touch, to take pleasure in one another's successes, and so on. Friends also have certain positively or negatively valenced *attitudes* towards one another in characteristic sorts of circumstances – gratitude, forgiveness, disappointment, and so on. Friends also tend to *expect* of one another (in a normative sense of “expect”) – absent exculpating conditions – that they will have these sorts of intentions and attitudes at the right times and in the right ways.

What is distinctive of our epistemic relations is how their characteristic intentions, expectations, and attitudes are oriented around our epistemic agency (Boult 2021; forthcoming). Epistemic relations are characterized by a kind of concern for the cultivation and utilization of epistemic agency, which is manifested in a unique set of intentions, expectations, and attitudes. Sometimes these arise out of contingent shared projects, such as being members of a scientific research team, or more mundanely, jointly searching for a lost set of keys. But at the most basic level, there is a sense in which epistemic agents, as such, stand in epistemic relations with others simply in virtue of being creatures who are capable of concern for the epistemic justifiability of their inquiries and attitudes.¹⁵ This can be seen in the way that, unless we have good reason not to, and within certain contextually determined domains, we tend to *epistemically trust* the word of others – even strangers – and expect that they will do the same, unless they have good reason not to.¹⁶ When I ask

¹⁵ I'm deeply indebted to Scanlon's (1998) work on contractualism in my development of aspects of the good relations view.

¹⁶ Epistemically trusting someone for information entails, minimally, that one has confidence that the person is epistemically reliable, and that one is willing to rely on them as such. Other things being equal, it also entails that one judges that the person will do their best in acquiring and sharing the

you for the next train time, other things being equal, and absent any defeating evidence regarding your reliability, I will take you at your word. In taking you at your word, I minimally reveal that I predictively expect that your answer will be worth heeding. But also, were I to discover you gave me egregiously inaccurate information, I might feel frustrated, perhaps even mildly resentful of your inaccurate assertion – unless I think you’ve got a good excuse. I might criticize or blame you, perhaps in a distinctively epistemic way (Boult 2021; Brown 2020; Piovarchy 2021; Schmidt 2024; cf. Kauppinen 2018). Many hold that the appropriateness of blame entails a kind of responsibility, suggesting my reaction entails that I hold you to some degree responsible for an accurate assertion. In other words, I have a certain degree of *normative* expectation that you will have a certain degree of competence relevant to my inquiry, and will give me an accurate answer. Meanwhile, in giving me the next train time, it’s also broadly familiar that, other things being equal, you will *intend* to do your best to provide me with accurate information, and you will expect – in a normative sense – that I will take your word seriously. Again, consider how you’d react if I simply rebuffed your assertion without explanation.

4. Epistemic Self-Relations

In addition to epistemic relations with others, people stand in epistemic relations with themselves. This follows straightforwardly from the schema so far: each of us stands in a relation of epistemic dependence with ourselves, and it is characterized by structurally similar intentions, expectations, and attitudes. Just consider the extent to which you need to epistemically trust yourself in order to get around in the world.

In his work on the social foundations of epistemic normativity, Sandy Goldberg (2018) observes that agents have expectations of themselves in virtue of standing in relations with later *time slices* of themselves. You’ve got a big deadline at the end of this week, and much else besides. You can’t do everything at once so you intend to work on the project starting Wednesday. In forming this intention, you

relevant epistemic goods, and will let one know when they are *not* reliable on some subject matter (Dormandy 2020; Woodard 2023).

normatively *expect* that your Wednesday-self will have the competences necessary to complete the project. There is an epistemic dimension to this: you expect you'll be able to form justified beliefs about what you're doing, reason well, conduct inquiry appropriately, and so on. In forming the intention to start work on Wednesday, you normatively expect yourself to be epistemically trustworthy. The good relations view takes this normative expectation to be a constitutive element of good epistemic relations.

In addition to time slices, we can also appeal to dimensions of the self at a single time slice. To take a practical case, suppose you're battling an addiction, or going on a diet. These projects highlight a kind of dividedness that can occur amongst a single individual's total intentions and desires (*akrasia*). To give this a Frankfurtian spin, the more fully one's first- and second-order desires "mesh" – the more one endorses one's effective desires at a second-order level – the better one's relations with oneself are. If there is such thing as *epistemic akrasia*, or the possibility of greater or lesser "meshing" between one's first- and second-order epistemic attitudes (beliefs, and beliefs about what one ought to believe, for example), a parallel case might be made for the idea of better and worse epistemic relations with oneself.¹⁷ Perhaps the most plausible way to frame this is again in terms of a relation of epistemic self-trust (cf. El Kassar 2019; Jones 2012).

Epistemic self-relations are among the most important (epistemic) relations we have. We rely on ourselves on a continual basis, including in – and as a background condition on – our practice of relying on others. If it turns out our epistemic self-relations are integral to, necessary for, or perhaps even partly constitutive of, our epistemic relations with others, this highlights an additional level of priority or importance we can attribute to our epistemic self-relations. These points will become important shortly.¹⁸

¹⁷ See Scanlon (1998, 33-38) for a compelling defense that treats epistemic *akrasia* as entirely parallel with practical *akrasia*. See Hurley (1989, 130-135) for an early argument against the possibility of epistemic *akrasia*.

¹⁸ I do not have space to develop the point here, but in my view epistemic relations can also obtain between *groups*, and between groups and individuals. In Boulton (MS), I flesh out these points and develop an account of how impairments to our epistemic relations at the group level can illuminate important social and political issues, such as epistemic decolonization.

5. Valuing Good Epistemic Relations

I've illustrated two ways epistemic agents can have reciprocal (and intrapersonal) sets of intentions, expectations and attitudes that are distinctively oriented around our epistemic agency. These intentions, expectations, and attitudes are at once sufficiently rooted in the epistemic domain to underpin an *epistemic* kind of relation, and sufficiently resemble the sort of intentions, expectations, and attitudes that make up other paradigmatic interpersonal relations that they can ground claims about final value. I turn now to a defense of this latter claim.

Many interpersonal relations seem to have final value, or in other words, to be valuable as ends in themselves or for their own sake. Friendships and other kinds of loving relationships are a case in point (Frankena 1973, 87; Moore 1903, 188; Scanlon 1998, 88; Stocker 1981). Might it be that, just in virtue of sufficiently resembling these interpersonal relations, our epistemic relations can have final value? I think this would be too quick.

Some relationships seem capable of giving us *reasons* to do or feel certain things, simply in virtue of being the kinds of relationships they are (Scanlon 1998, Ch.2; Stocker 1981, 752). It would seem odd to say you've got a reason to visit your mother in the hospital because she's a nice, thoughtful, and intelligent person (Kolodny 2003, 139). Other things being equal, the more natural thing to say is that you've got a reason to do so because she's *your mother*. Replace your mother with some random but equally nice, thoughtful, and intelligent person in the hospital. Other things being equal, you'd no longer have a reason to go to the hospital. The ability to give us reasons in their own right may suggest that certain relationships have final value, and that they can be properly valued for their own sake. The trouble with this observation is that it may highlight an important *difference* between some kinds of relationships – such as loving ones, or personal relationships more generally – and our epistemic relations. It may seem implausible, for example, that the sheer fact of my standing in an epistemic relation with someone with whom I otherwise have no relations whatsoever, can give me a reason to do anything.

But it is important to be clear about the kinds of reasons personal relationships give us. We said that the sheer fact it's your *mother* can give you a reason to visit her in the hospital. Here, your reason to go to the hospital is a token instance of a more general type of reason: a reason you have to live up to certain expectations constitutive of the normative ideal of that relationship. Stated like this, it becomes more plausible that epistemic relations, too, can give us reasons to do things. The fact that you stand in an epistemic relation with S—regardless of who S is, or whatever properties S has—gives you a reason to live up to the expectations partly constitutive of the normative ideal of that relation. That said, it's natural to wonder whether, even if this is true, it is simply because of the instrumental value of standing in good epistemic relations. If so, whatever reasons our epistemic relations give us, it need not be any indication that good epistemic relations have final value.

I propose that the fact that you stand in an epistemic relation with S—regardless of who S is, or whatever properties S has—gives you a reason to live up to the expectations partly constitutive of the normative ideal of that relation, not necessarily because of any instrumental value that accrues in doing so, but rather because doing so is a way of *recognizing* the epistemic agency of others. Having a concern for the epistemic justifiability of one's attitudes and inquiries just is, in part, a way of recognizing and properly responding to epistemic agency. To my mind, there is an intuitive sense in which recognizing and properly responding to epistemic agency is good for its own sake, perhaps not unlike how living up to the normative ideal of your relationship with your friends can be good for its own sake. Regardless of whether doing so is useful, a world in which epistemic agents recognize and properly respond to one another's properly functioning epistemic agency is better than a world in which they do not. If that is right, then good epistemic relations may have final value. A complete defense of these claims would take more space than I have here. But it is enough for now to establish the following conditional: if good epistemic relations have final value, a powerfully ecumenical and independently well-motivated response to the source question is just a short

step away. Even if I cannot now fully defend the antecedent, I will be satisfied if I can demonstrate the promise of exploring this conditional.¹⁹

According to a very broad distinction between ways of valuing, one way to value something is by promoting it, another is by *respecting* it (Parfit 2011, 236-237; Scanlon 1998, Ch 2; Stocker 1981; Sylvan 2020). Promoting requires doing things or having attitudes that conduce to the existence of the item of value. What is respect? I follow Kurt Sylvan (2020) in distinguishing between two kinds of respect: esteem and “recognition respect” (Darwall 1977). Esteem is the kind of respect at issue when we express admiration, as when I say “I respect Barack Obama”. Recognition respect can be captured as a kind of responsiveness to reasons (where “V” is an item of value):

- A person *S* *weakly respects* V in Φ -ing iff Φ -ing is favored by *S*’s merely subjective V-related reasons, and *S* Φ s for these subjective reasons.
- A person *S* *strongly respects* V in Φ -ing iff Φ -ing is favored by *S*’s evidence-relative V-related reasons, and *S* Φ s for these evidence-relative reasons.

As Sylvan notes, “There is something good from the epistemic point of view about manifesting each form of respect, but the second is better than the first” (2020, 17). Weak respect is a kind of *conscientiousness*. But one can be irrationally conscientious, or massively mistaken about what one’s evidence really supports. In what follows, I

¹⁹ One question that might seem pressing is: finally valuable *for who*? First, I am open to the possibility that good epistemic relations are finally valuable simpliciter. Not all things need to be valuable *for* someone (Way 2013, 7). That said, since epistemic relations at the most basic level obtain between agents capable of concern for the epistemic justifiability of their attitudes and inquiries, I also think it’s plausible that, in the most basic sense, good epistemic relations are finally valuable for anyone capable of concern for the epistemic justifiability of their attitudes and inquiries. It’s also worth emphasizing, as I noted earlier, that in addition to the most basic level, we stand in epistemic relations with others and ourselves in virtue of contingent projects, such as being members of a research team. Dimensions of our epistemic relations that arise out of such projects may only be good for particular individuals, such as the members of a research team. In this way, the good relations view can accommodate the plausible idea that we may not all have the same reason to comply with the same epistemic norms, nor to stand in the same types of epistemic relations with all people. See Boulton (forthcoming, Ch. 6) for a detailed development of some of these ideas. Thanks to Antti Kauppinen for helpful comments here.

am more interested in strong respect. One successful and particularly interesting way of valuing something is by strongly respecting it.

What sorts of things are properly valued through respect rather than by performing actions or having attitudes aimed at promoting their existence? Human dignity and autonomy may be good examples. Certain kinds of interpersonal relationships, such as friendship, also seem like candidates (Parfit 2011; Scanlon 1998, 89; Stocker 1981). Consider: even if one could somehow bring about the existence of three new friendships by betraying the trust of one friend, betraying the trust of that one friend would not be a way of properly valuing friendship. So, merely promoting the existence of good friendships does not exhaust the ways we can properly value friendships. Sometimes properly valuing them requires respecting them in certain ways.

As a kind of interpersonal relation with final value, my view is that properly valuing good epistemic relations consists not only in promoting their existence but also respecting them. Let V^{er} be the value of good epistemic relations. The V^{er} -related reasons are the reasons that favour having the kinds of intentions, expectations, and attitudes constitutive of good epistemic relations. If complying with epistemic norm n is favoured by one's evidence-relative V^{er} -related reasons, and one complies for those reasons, then complying with norm n can be a way of respecting V^{er} . Why think respect and not just promotion is a way of properly valuing good epistemic relations? Consider: even if one could somehow bring three new good epistemic relations into existence by violating the expectations of just one, this would not be a way of properly valuing good epistemic relations. So, merely promoting the existence of good epistemic relations does not exhaust the ways we properly value epistemic relations. Sometimes properly valuing them requires respect.

Good epistemic relations are significant because, as a kind of inter- and intrapersonal relation with final value, it is independently plausible that properly valuing them includes not just promoting them but respecting them. One useful thing about this is that it seems less controversial and abstract than, for example, Sylvan's claim that "the fundamental epistemic value" (on his view, accuracy) calls for respect in the first instance (2020, 10). For my own part, I do not find it obvious

that the proper way of valuing accuracy is by respecting it. But I don't need to commit to this point. My view merely trades on the idea that good epistemic relations call for respect.

I want to emphasise that final value is not incompatible with instrumental value. When I say that a world in which I have friends is better, for its own sake, than one in which I do not, I am not denying that having friends can be useful (cf. Sylvan 2020, 8). I am only saying that their usefulness does not exhaust an account of their value, and how we properly value them. The same is true about good epistemic relations. As I will argue, the connection between our epistemic relations' usefulness and their final value is key to the ecumenical power of the good relations view.

6. Problems for Social Instrumentalism and How Good Epistemic Relations Can Help

According to the good relations view, members of epistemic communities stand in epistemic relations that can be better or worse. Epistemic norm compliance is a necessary condition on participation in good epistemic relations, and good epistemic relations are i) integral to individuals' and their communities' ability to flourish, and ii) something properly respected for their own sake. So, we have reason to comply with epistemic norms. The view clearly shares key elements with existing forms of social instrumentalism, such as those inspired by Chrisman (2022), and directly developed by Hannon and Woodard (forthcoming) (Section 2).²⁰ While agreeing that we tend to rely on one another to comply with epistemic norms as a way of promoting the kind of *coordination* and *knowledge acquisition* that seems integral to flourishing human life, the good relations view adds that this sort of reliance – underpinned as it is by reciprocal sets of intentions, expectations, and attitudes that are oriented around a concern for the epistemic justifiability of our attitudes and inquiries – amounts to a distinctive interpersonal relation that has final value.

As I argue next, this opens up novel space to address some central challenges for social instrumentalism. The good relations view has resources for dealing with important worries about the grip of epistemic norms in the absence of social

²⁰ The view also has close affinities with Goldberg's 2018 approach to the source question.

frameworks, and (empirically based) considerations about epistemically flourishing social frameworks in the absence of epistemic norm compliance.²¹ To be clear, my aim isn't to conclusively argue there are no alternative avenues for social instrumentalists to deal with these central challenges (though I outline some relevant concerns). It's to argue that the good relations view offers a comparatively simple and elegant way of doing so.

6.1 Norm Compliance Without Social Frameworks

Suppose someone's epistemic conduct has no chance of impacting (or being impacted by) the epistemic community, whether by choice – say, by deciding to live alone in a remote part of the Yukon – or unintentionally – say, by being depressed to such a degree that one is incapable of participating in the community. Don't epistemic norms apply to these agents? How can a view claiming that people should comply with epistemic norms because doing so enables the success of social epistemic practices that are integral to the achievement of both individual and collective goals explain this?

Chrisman directly addresses the problem of social isolation, pointing out that a more individualistic story about epistemic normativity can account for it (2022, 138). For such agents, even though they are isolated from the social practice, they nevertheless have their own goals, and complying with epistemic norms is a means to promoting those goals, so an individualistic instrumentalism remains in play (cf. Dyke 2021). For Chrisman, social instrumentalism isn't the *whole* story, it's just a story about how individualistic instrumentalism is incomplete.

I'm not convinced this gets at the heart of the issue. Shunting the burden back onto individualist instrumentalism in social isolation cases leads back to the too few reasons problem. We can imagine the relevant kinds of disconnect between epistemic norms and individual goal promotion in agents who are causally isolated from the social practice. Socially isolated agents are not immune from the too few reasons problem.

²¹ See also fn. 26 for an example of how lessons from these issues can be more widely applied.

Hannon and Woodard don't directly engage with social isolation, but may offer insights nonetheless. In line with Section 2, they take the central issue for direct instrumentalism to be the too few reasons problem (calling it "the central problem"). To deal with agents for whom complying with epistemic norms would frustrate or fail to promote any goals, Hannon and Woodard point out it is nevertheless legitimate for other members of the epistemic community to *criticize* and *epistemically blame* agents if they flout those norms. This is because doing so is a way of upholding the epistemic norms of the community, thereby ensuring the possibility of the achievement of many social (and individual) goals (Hannon and Woodard forthcoming). But they are careful to point out that none of this implies such agents have a genuine *reason* to comply with epistemic norms. The appropriateness of criticism and blame merely indicates that epistemic norms "apply" to these agents irrespectively of their own personal goals. In effect, Hannon and Woodard draw a distinction between two kinds of categoricity – "strong" (epistemic norms give us *reasons*, regardless of our goals) and "weak" (epistemic norms *apply* to us, regardless of our goals) – and argue that epistemic norms are only weakly categorical. They endorse an error theory about the appearance of the "authority" of epistemic normativity (forthcoming, 17).

Perhaps even if isolated agents (including the kind at issue in the too few reasons problem) lack reasons to comply with epistemic norms, it is nevertheless appropriate to epistemically criticize them for bad epistemic conduct. Epistemic norms may still in an important sense "apply" to these agents. And perhaps this is all we really need in order to make sense of the idea that isolated agents (including the kind at issue in the too few reasons problem) should, in *some* sense, comply with epistemic norms.

The trouble with this argument is it's unclear why it would be appropriate to so much as epistemically *criticize* or blame socially isolated agents, at least by Hannon and Woodard's lights. On their view, epistemic criticism is typically appropriate because it is a central means for members of the epistemic community to *enforce* epistemic norms, the general successful enforcement of which in turn allows us to achieve many of our individual and collective goals. To quote Hannon and

Woodard, “epistemic norms are *in force* because they are *enforced*” (forthcoming, 14). In isolation cases, this story seems to make little sense. These agents are, by definition, ones whose epistemic conduct makes no difference to the broader epistemic community (and vice versa). So, if the function of epistemic criticism is to ensure that agents who might otherwise make a detrimental difference to the epistemic community refrain from doing so, epistemic criticism would seem out of place here – if only because members of the epistemic community would be better off placing their efforts elsewhere. In what sense, then, are these agents epistemically criticizable?

A natural answer is that they have flouted their reasons to comply with epistemic norms. Regardless of whether they have any capacity to make an impact on the community, and regardless of whether complying with epistemic norms would promote their goals, if one has a reason to comply with epistemic norms, and flouts that reason, perhaps they are epistemically criticizable. While natural, this answer is unavailable to Hannon and Woodard. They would deny that socially isolated agents (of the kind at issue in the too few reasons problem) have reasons to comply with epistemic norms (even if those norms “apply” to such agents).

Hannon and Woodard might simply embrace the idea that such agents aren’t epistemically criticizable, and so perhaps epistemic norms do not even apply to them. But this leads to at least two other points. First, embracing this stance goes against the spirit of their original error theory. They invoke the idea of “norm applicability” as a central mechanism for appeasing our intuitions about the categoricity of epistemic normativity. If they reject those intuitions tout court in isolation cases, why bother appeasing them in the first place? Second, embracing the idea that these agents aren’t criticizable doesn’t get at the heart of what Hannon and Woodard seem unable to satisfyingly account for here – namely, the intuitiveness of the idea that even socially isolated agents (including the kind at issue in the too few reasons problem) have reason to comply with epistemic norms.²²

²² The plausibility of this point may require restriction to specific kinds of epistemic norms. For sake of argument, focus on the norm enjoining us to proportion our beliefs to the evidence.

How can the good relations view do better? This may seem unclear. After all, socially isolated agents don't stand in epistemic relations with others. So how can a view that explains our reason to comply with epistemic norms through the role that epistemic norm compliance plays in good epistemic relations work here? First, recall from Section 4, I argued that one of the most important epistemic relations we stand in is the one with ourselves. This is true in virtue of how we rely on later time-slices of ourselves, and also in virtue of relations of epistemic trust that can take place between orders of epistemic attitudes we hold synchronically. Taking inspiration from Goldberg (2018, Chapter 7.4), we can press on this observation now. Just like our epistemic relations with others, good epistemic self-relations are both instrumentally and finally valuable. Complying with epistemic norms, and maintaining epistemic self-trustworthiness, is a constitutive element of maintaining good epistemic self-relations. So, since good epistemic self-relations have both instrumental and final value, socially isolated agents have reason to comply with epistemic norms, even when doing so does not directly promote any of their individual goals.

Perhaps a structurally similar move is available to my competitors. Chrisman might suggest that socially isolated agents have an interest in maintaining a good epistemic reputation with *themselves*. Hannon and Woodard might argue that socially isolated agents are proper subjects of their own epistemic criticism (say, "epistemic guilt"). Perhaps in virtue of the relevant promotion and avoidance relations, each of these accounts can equally readily explain why epistemic norms apply to socially isolated agents after all.

To my mind, neither of these suggestions is plausible. The idea of maintaining an epistemic reputation with oneself strikes me as an idle theoretical wheel: why would we need to maintain a reputation with ourselves if we already have more direct (albeit imperfect) routes to information about our own reliability and trustworthiness? Reputation maintenance seems to have an inherently *interpersonal* function. Meanwhile, the idea of epistemic self-criticism is surely coherent (and perhaps important), but recall the *function* of epistemic criticism on Hannon and Woodard's view. On their view, we epistemically criticize each other as a means of

ensuring proper compliance with epistemic norms across the community, in light of how epistemic norm compliance promotes knowledge acquisition and coordination. It's unclear to me whether epistemic *self*-criticism can play this kind of productive role in our epistemic lives, especially on the assumption of social isolation. For a start, it seems inherently more difficult for criticism to function as a productive corrective when evaluations and objects of evaluation issue *exclusively* from the same perspective (i.e. the socially isolated individual's). Still, I don't deny that epistemic self-criticism may have a role to play in the solo case, perhaps one resembling the role interpersonal criticism plays across epistemic communities. In this respect, the good relations view and Hannon and Woodard's social instrumentalism may be on similar ground. So it's worth highlighting that the good relations view has additional resources in this context as well.

Recall the distinction between two ways of valuing: promotion and respect. It may of course seem hopeless to try and explain why socially isolated agents have reason to comply with epistemic norms by appealing to the role this plays in maintaining good epistemic relations with others, since socially isolated agents don't stand in any such relations. But perhaps this only seems odd if we think of the explanatory connection as a promotion relation – that is, in terms of the idea that compliance with epistemic norms is a way of valuing good epistemic relations *by promoting their existence*. As I have argued, this is not the only way we properly value things. And indeed, it is implausible that the proper way of valuing good epistemic relations is exhausted by promoting their existence. While socially isolated agents cannot do anything to promote the existence of epistemic relations with others, they can nevertheless respect the final value of good epistemic relations. Compare how people find it perfectly natural to respect someone who is long dead, and possibly carry out elaborate actions in their efforts to do so. We regularly respect things even in the absence of the possibility of promoting their existence. In my view, compliance with epistemic norms – even if it does nothing to promote one's goals, or to promote the existence of epistemic relations with others – is a way of respecting the final value of good epistemic relations. This is an additional way of explaining part of the grip of epistemic norms in such cases.

6.2 Social Frameworks Without Norm Compliance

Another central worry for social instrumentalism is that it is ultimately an empirical claim whether individuals' compliance with epistemic norms is really as important for human flourishing as social instrumentalism presupposes (Smart 2018; Bland 2022; Henrich 2016). Recent work by Neil Levy and Mark Alfano (2020) brings this worry into sharp relief.²³

Drawing on a range of empirical work, Levy and Alfano argue that, over time, many instances of human flourishing (including basic survival) are only made possible through the intellectually *vicious* conduct of individual members of a group or society. They use "cumulative culture" as an example of how solutions to problems involving causally opaque mechanisms can require a kind of blind deference to tradition and close-mindedness (cf. Henrich 2016; Bland 2022, 59). Groups who perpetuate cultural practices through unreflective imitation, deference, and false beliefs about causal processes, often end up with highly adaptive solutions to causally opaque challenges in their environment, solutions that would be nearly impossible to come up with through an individual or group's coordinated exemplary adherence to the epistemic norms enshrined in mainstream analytic epistemology (Levy and Alfano 2020, 895). Their central examples include elaborate preparation rituals for (opaquely) toxic plants (2020, 894; cf. Bland 2022, 59; Henrich 2016, 30), and the consultation of auguries in the planning of hunt routes (Levy and Alfano 2020, 904). According to Levy and Alfano, these are examples of practices that have proved adaptive for certain groups precisely because of certain false beliefs²⁴ and blind deference in the respective groups' adherence to them. The thought is that rather than requiring compliance with epistemic norms, empirical evidence suggests that human flourishing often comes about in precisely the

²³ I will focus on Levy and Alfano, but I take my point to apply to a wide range of views similar to Levy and Alfano's. See Bland (2022) for a helpful overview of such views.

²⁴ In the case of auguries in planning hunting routes, there is evidence (Henrich 2017) that this practice increased hunting yields because it effectively *randomized* behaviour, whereas practitioners believed the auguries were divine messages (Levy and Alfano 2020, 904).

opposite way, through close-mindedness, lack of curiosity, and excessive deference to others.

In response to this argument, some have argued that the empirical evidence suggests that individuals receiving testimony in the transmission of cumulative cultural knowledge deserve more epistemic credit than Levy and Alfano conclude (Sperber et al. 2010; Mercier 2017). Others have argued that the epistemic credit at issue in the transmission and acquisition of cumulative culture is best understood as *distributed* across individuals (Palermos 2016). Social instrumentalists might avail themselves of these strategies, using them to defend the role that epistemic norm compliance plays (perhaps despite appearances, or in a distributed way) in human flourishing, even in Levy and Alfano's putatively problematic cases of cumulative culture. One significant concern here, however, is that getting clear on whether these strategies work is a massively complex empirical issue, something unlikely to clear up any time soon. It is well worth advertising, then, that the good relations view has an entirely different resource in this context, one that allows us to sidestep empirical controversy.

According to the good relations view, even if things like close-mindedness and blind deference are integral to cumulative culture and in turn underpin certain highly adaptive forms of co-ordination and knowledge acquisition for communities, the extent that they involve flouting epistemic norms²⁵ remains a way of falling short of good epistemic relations. Insofar as regularly flouting epistemic norms manifests a lack of concern for the *epistemic justifiability* of one's attitudes and inquiries to others, it constitutively involves a certain lack of recognition of one's own and others' epistemic agency. If I am right that mutual recognition of epistemic agency has final value – a kind of value that floats free of instrumental considerations – it follows that this is something of a loss. And so, despite the instrumental benefits of certain intellectual vices in certain contexts (assuming this is empirically well-supported), we still have reason to comply with epistemic norms, even in those very contexts. Perhaps these empirical observations show how the grip of epistemic

²⁵ Again, by “epistemic norms”, I mean any part of the epistemic domain that seems to have the normative grip that epistemologists are interested in.

normativity in some sense “loosens” in certain ways. But in my view, no matter what long-term adaptive benefits accrue to collectives through the intellectually vicious conduct of individuals, there remains a sense in which individuals have reason to comply with epistemic norms. As Levy and Alfano themselves write: “While we maintain that these supposed vices are knowledge-conducive much more often and more reliably than others have recognized, there is *something* to the identification of opposing dispositions as virtues” (Levy and Alfano 2020, 515 emphasis mine). The good relations view has a simple explanation of what this “something” is: insofar as intellectual virtues involve dispositions to comply with epistemic norms, they are a constitutive part of what it takes to respect and promote good epistemic relations. The good relations view can explain why agents engaged in epistemically vicious practices that facilitate cumulative culture do have reason to comply with epistemic norms, even if not an overriding reason.

Non-ideal epistemologists like Levy and Alfano often say that a central problem with traditional epistemology is that it’s too *individualistic* to appreciate social insights from empirical research. Interestingly, my concern in this section doesn’t stem from an individualistic epistemology. It stems from a deeply social one. This may suggest that a diagnosis of hyper-individualism misses the heart of the matter. One need not be hyper-individualistic to be interested in the grip of epistemic norms. One can be interested in that grip insofar as it arises out of the quality of our epistemic relations. Given that good epistemic relations ought to be respected, this explanation of the grip of epistemic norms is consistent with the sorts of empirical claims non-ideal epistemologists make about how we can best promote items of epistemic value.²⁶

7. Is the Good Relations View Circular?

²⁶ For some forms of social instrumentalism, a challenge related to the one in this section is an epistemic version of the free-rider problem—cases of individuals living flourishing human lives, or otherwise reaping the benefits of our epistemic practices, without regularly complying with epistemic norms (say, by being consummate bullshitters). The good relations view can invoke the idea that such agents miss out on the final value of good epistemic relations to account for these cases as well.

The good relations view aims to answer the question: why comply with epistemic norms? And its answer is, in part, because doing so is a constitutive element of maintaining good epistemic relations. When dealing with the further question, “what are good epistemic relations?”, it may seem as though the account’s answer ultimately boils down to: “relations between agents who comply with epistemic norms!” This may seem worryingly circular, or otherwise uninformative. In this section, I argue it is not.

To start, as I noted in the Introduction, there is an important difference between trying to convince a *skeptic* about epistemic normativity that they have reason to comply with epistemic norms, versus explaining to someone who is already convinced that epistemic norm compliance is something we have reason to do, *why* we have reason to do it. I concede that the story above may not convince a skeptic. But I’m also not particularly worried about the account’s anti-skeptical powers. For one thing, I’m not convinced such a skeptical stance can coherently be adopted. It seems that any legitimate challenge to the idea that we have reason to comply with epistemic norms would itself require – however implicitly – reliance on the idea that we have reason to comply with epistemic norms. What reason would we have to believe it’s a legitimate challenge otherwise? Regardless, even if the good relations view can’t satisfy a skeptic, it does explain to the person who’s already a fan of epistemic normativity why they have reason to comply with epistemic norms. It puts structure on our understanding of the role of epistemic norms in our lives.

There is a difference between a set of mutually supporting ideas and ideas that stand in a viciously circular relationship. The nexus of ideas at play in the good relations view is mutually supporting. Partly this is because good epistemic relations are not simply ones in which participants comply with epistemic norms. They are also ones in which people epistemically *trust* each other when they are epistemically trustworthy. Being epistemically trustworthy means reliably complying with epistemic norms. But it also means that, for any two members, A and B, concerning some proposition *p*, under the right circumstances, A knows that B may epistemically rely on them for information pertaining to *p*; that A is aware of their epistemic limitations regarding *p*; and in many cases, that A knows that B knows

that A knows that B may epistemically rely on them for information pertaining to p, and so on. Epistemically trusting a person with respect to p is standardly regarded as involving more than just knowing or believing that a person tends to comply with epistemic norms around the issue of whether p. Epistemic trustworthiness with respect to p is standardly regarded as more than just a proclivity to comply with epistemic norms around the issue of whether p (Faulkner 2011; Dormandy 2020; Woodard 2023).

As I put it earlier, good epistemic relations manifest a kind of recognition respect for epistemic agency. This is arguably what makes them finally valuable. In explaining our reason to comply with epistemic norms by appealing to the role doing so plays in promoting and respecting the final value of good epistemic relations, the good relations view makes use of mutually supporting theoretical ideas, but it is not viciously circular.²⁷

8. Ecumenical Power

In relying as heavily as I do on the final value of good epistemic relations, we might wonder whether my view is ultimately just a form of *anti-* or at least *non-*instrumentalism. In explaining our reasons to comply with epistemic norms in terms of the idea that doing so is a way of respecting the final value of good epistemic relations, my account circumvents any reference to goals or aims at a crucial

²⁷ Here is a different concern about the informativeness of the good relations view. Even if good epistemic relations have final value, and even if the account is not viciously circular, it remains a question whether the account illuminates why we *ought* to do anything at all. Just because something is valuable, it doesn't follow that you ought to do anything to promote or respect it. It would be good if I gave more to charity. But it would also be good if I saved more for my daughter's education. Perhaps I ought to do the latter at the expense of the former. Even if complying with epistemic norms is a way of promoting and respecting the final value of good epistemic relations, the worry goes, this doesn't explain why we ought to comply with epistemic norms. In response, note that answering the source question does not require explaining what it is in virtue of which we ought to comply with epistemic norms once all the facts are in. As I said at the outset, answering the source question requires explaining the force or grip of epistemic norms – where, again, “epistemic norms” is a catch-all term referring to anything in the epistemic domain that seems to have the kind of grip or force that philosophers have been interested in. We can explain this grip, allowing that it may sometimes loosen (as I put it in the context of cumulative culture in Section 6.2), even if doing so falls short of explaining what anyone really ought to do at any given moment in time. A complete account of the weighting issues this may give rise to is beyond the scope of this paper. Thanks to Antti Kauppinen for raising this issue.

juncture. I'm presenting the good relations view as an "ecumenical" approach to the source question. But have I not simply abandoned instrumentalism altogether?

The ecumenical power of the account lies in the structural and explanatory connection between its claims about final value and the social instrumentalist framework in which they are embedded. The mediating factor is our epistemic relations. While good epistemic relations clearly have instrumental value much of the time, we must acknowledge that there are possible cases – as discussed above – in which they do not, or where individuals can just as well achieve their goals without them. In those sorts of cases, however, the independently plausible idea that good epistemic (self-)relations have final value explains our reason to comply with epistemic norms. Doing so is just a distinctive way of valuing this item of final value. The tight connection between the instrumentalist and non-instrumentalist aspects of the account sets it importantly apart from existing *anti*-instrumentalist views, and thereby provides one way of highlighting its comparative ecumenical power. Let me explain.

Consider that nearly all anti-instrumentalists maintain that complying with epistemic norms is instrumentally valuable much of the time. This observation is a platitude (Alston 2005; Dyke 2021; Willoughby 2022). One difference between anti-instrumentalist approaches to the source question and the good relations view is that, on standard anti-instrumentalist frameworks, the instrumental value of epistemic norm compliance seems *coincidentally correlated* with our reasons to comply with epistemic norms. Christopher Cowie (2014) articulates the point in an "argument from coincidence", presenting it as a deep challenge for anti-instrumentalism. As he puts it:²⁸

The intrinsicist is committed to the existence of both a practical value and a quite independent brutely epistemic value in believing in accordance with one's evidence. This is a very striking coincidence: there happens to be a brutely epistemic value in forming beliefs in a fashion that quite

²⁸ Cowie uses the term "intrinsicist" for views I am calling "anti-instrumentalist".

independently is already of great practical utility (and hence, value). Whilst accidents – even fortunate ones – do happen, this is a striking one indeed (2014, 4008).

The idea that agents have reason to comply with epistemic norms because, say, doing so is a constitutive element of being a believer, or because of facts about the etiological function of belief, or because this is just a “brute normative fact”, all seem to stand in stark isolation from direct claims about the instrumental value of complying with epistemic norms. They seem like *separate* observations about epistemic norms. As such, according to Cowie, they are coincidentally correlated with the fact that epistemic norm compliance tends to be instrumentally valuable. And according to Cowie, this seems like *too* much of a coincidence.

Cowie considers potential responses to the argument from coincidence. The most compelling he considers are i) the idea that there is a *common factor* explaining both the instrumental and epistemic value of complying with epistemic norms, and ii) the idea that there is a genealogical connection linking these two data-points in a non-coincidental way. Starting with the common factor approach, perhaps the most obvious candidate is *truth-conduciveness*. The thought would be that compliance with epistemic norms is at once instrumentally valuable because they are truth-conducive, and epistemically valuable because they are truth-conducive. This may seem promising, but I agree with Cowie that it just pushes the striking coincidence further back: either the *epistemic* value of truth-conduciveness derives from its usefulness – in which case we’ve circled back to instrumentalism – or it does not; but then it remains a striking coincidence that truth-conduciveness is both instrumentally and epistemically valuable.²⁹

Regardless of whether the argument from coincidence is a fatal problem for anti-instrumentalism, my aim here is to advertise the converse situation for the good relations view. One way of framing my approach is as a different kind of common factor strategy. The idea that there is a “structural connection” between my claims

²⁹ On genealogical or “selection-based” responses to the argument from coincidence (the kind etiological functionalists might pursue), see Cowie (2014, 4009).

about instrumental and final value can be put in terms of how good epistemic relations are a *common factor* between these claims, eliminating any appearance of coincidence. Good epistemic relations are instrumentally valuable. But it's also true that, in virtue of being the sorts of things they are (manifestations of a recognition of epistemic agency), they have final value worth respecting in its own right. Cowie motivates his point that a truth-conduciveness-based common factor strategy merely pushes the bump in the rug by relying on the rhetorical question: if not the ultimate usefulness of truth-conduciveness, what *else* could explain the importance of truth-conduciveness? The good relations view has a ready answer to the question of what else besides instrumental value explains our reasons to comply with epistemic norms. Good epistemic relations have final value. And compliance with epistemic norms is a way of properly valuing this item of final value. This is not an *ad hoc* addition to an otherwise instrumentalist answer to the source question. It is an appeal to the common factor between the account's dual claims about value: epistemic relations.

A standard complaint about anti-instrumentalism is that it simply doesn't *explain* our reason to comply with epistemic norms. Indeed, some versions of the view explicitly treat epistemic normativity as explanatorily basic, or a brute normative fact (Kelly 2003; Worsnip 2022). Even if the good relations view is ecumenical in the way I've just argued, is appealing to the final value of good epistemic relations any more theoretically satisfying than anti-instrumentalist explanations of the source of epistemic normativity?

The final value of good epistemic relations is appealing, in comparison with, say, brute normative facts, because the final value of good epistemic relations is a species of something we already routinely regard as having final value – namely, good human relations more generally. There may be something explanatorily basic about the claim that good epistemic relations have final value. But this claim about final value is a species of a much more general claim that already enjoys a great deal of plausibility and support – amongst instrumentalists and anti-instrumentalists about epistemic normativity alike.

One might object that, even if human relations more generally sometimes have final value, *epistemic* relations are different enough from those kinds of interpersonal relations – for example, more personal ones, such as friendship – that this species-genus defence of the framework is spurious. I acknowledge that articulating the notion of “epistemic relations” in the way I have done reveals competing sources of pressure: on the one hand, we want to articulate something distinctively epistemic (and not, for example, some kind of relation having nothing to do with epistemic norm compliance); on the other, we want to articulate something that plausibly has final value. These two desiderata can easily pull in opposing directions. In response, I submit that I have already discharged the burden of reconciling these competing sources of pressure. Good epistemic relations *are* both distinctively epistemic and finally valuable. Indeed, I argued that good epistemic relations have final value not merely by appealing to the fact that I call them “interpersonal relations”. But rather by appealing to properties they have which seem capable of underpinning claims about final value – namely, the fact that good epistemic relations involve a kind of mutual recognition and proper response to epistemic agency.

I’ll close by briefly noting that we gain additional reason to take the theoretical usefulness of epistemic relations seriously from some perhaps surprisingly disparate areas of epistemology. Consider: epistemic relations – or “epistemic relationships” as they’re sometimes called – already do important work in making sense of a wide range epistemic *harms* (such as epistemic exploitation, gaslighting, epistemic injustice by credibility deficit, and epistemic injustice by credibility surplus), and the nature and norms of epistemic *blame* (our practice of criticizing one another for epistemic failings). Roughly, a wide range of epistemic harms are helpfully unified as a kind of *impairment* to our epistemic relations (Boult forthcoming, Ch 5). And epistemic blame is helpfully understood as a kind of *modification* to those relations (Boult 2021; Greco forthcoming; Flores and Woodard 2023; Schmidt 2024). If the source of epistemic normativity is our epistemic relations, perhaps it should come as no surprise that our way of responding to one another for not properly respecting epistemic normativity is by modifying those relations.

The broader usefulness of epistemic relations, and the corresponding coherence it brings to our theorizing, is further reason to take seriously the role I have given epistemic relations in answering the source question. Elaborating on this broader coherence is a project for another time. But I expect it will be an important part of a full defense of the relational foundations of epistemic normativity.

Bibliography

- Alston, W. 2005. *Beyond justification*. Ithaca: Cornell University Press.
- Bird, A. 2007. Justified judging. *Philosophy and Phenomenological Research*, 74(1), 81–110.
- Bland, S. 2022. In defence of epistemic vices. *Synthese*, 200(59).
- Boult, C. 2021. There is a distinctively epistemic kind of blame. *Philosophy and Phenomenological Research*, 103 (3):518-534.
- _____. forthcoming. *Epistemic Blame: The Nature and Norms of Epistemic Relationships*. Oxford: OUP.
- _____. MS. Epistemic reparations and our relations as knowers.
- Brown, J. 2020. What is epistemic blame? *Noûs*, 54 (2):389-407.
- Buckley, D. 2021. Varieties of epistemic instrumentalism. *Synthese* 198 (10): 9293–9313.
- Burge, T. 2003. Perceptual entitlement. *Philosophy and Phenomenological Research*, 67(3), 503–548.
- Côté -Bouchard, C. 2015. Epistemic instrumentalism and the too few reasons objection. *International Journal of Philosophical Studies*, 23(3): 337–355.
- _____. 2016. Can the aim of belief ground epistemic normativity? *Philosophical Studies* 173 (12): 3181-3198. 2016.
- Carter, A. 2022. Epistemic normativity is not independent of our goals. In Blake Roeber, Matthias Steup, Ernest Sosa & John Turri (eds.), *Contemporary Debates in Epistemology (3rd ed.)*. Wiley-Blackwell.
- Chrisman, M. 2022. *Belief, Agency, and Knowledge*. Oxford: OUP.
- Cowie, C. 2014. In defense of instrumentalism about epistemic normativity. *Synthese*, 191: 4003-4017.
- Darwall, S. 1977. Two kinds of respect. *Ethics*, 88(1):36-49.
- Dogramaci, S. 2012. Reverse engineering epistemic evaluations. *Philosophy and Phenomenological Research*, 84 (3): 513-530.
- Dormandy, K. 2020. *Trust in Epistemology*. New York: Routledge.
- Dyke, M. 2021. Could our epistemic reasons be collective practical reasons? *Noûs*, 55 (4): 842-862. 2021.
- El Kassar, N. 2019. The place of intellectual self-trust in theories of epistemic advantages. *Journal of Social Philosophy* 51 (1):7-26.
- Flores, C. and Woodard, E. 2023. Epistemic norms on evidence gathering. *Philosophical Studies* 180 (9):2547-2571.
- Foley, R. 1987. *The Theory of Epistemic Rationality*. Harvard University Press
- Frankena, W. 1973. *Ethics*. Englewood Cliffs, N.J.: Prentice-Hall

- Goldberg, S. 2018. *To the Best of Our Knowledge*. Oxford: OUP.
- Graham, P. 2015. Epistemic normativity and social norms. In David K. Henderson & John Greco (eds.), *Epistemic Evaluation: Purposeful Epistemology*. Oxford University Press. pp. 247-273.
- Greco, D. Forthcoming. On the very idea of an epistemic dilemma. In N. Hughes (Ed.), *Essays on Epistemic Dilemmas*, Oxford: OUP.
- Grimm, S. 2009. Epistemic normativity. In A. Millar, A. Haddock, & D. Pritchard (Eds.), *Epistemic value*. New York: Oxford University Press.
- Hannon, M. and Woodard, E. Forthcoming. The construction of epistemic normativity. *Philosophical Issues*.
- Henderson, D. 2020. Are epistemic norms fundamentally social norms? *Episteme* 17 (3):281-300.
- Henrich, J. 2016. *The Secret of Our Success: How Culture is Driving Human Evolution, Domesticating our Species, and Making us Smarter*. Princeton: Princeton University Press
- Hurley, S. 1989. *Natural Reasons*. Oxford: OUP. pp.130-135.
- Jones, K. 2012. The politics of intellectual self-trust. *Social Epistemology*, 12(2): 237-251.
- Kauppinen, A. 2018. Epistemic normativity and epistemic accountability. *Philosophers' Imprint*, 18.
- Kelly, T. 2003. Epistemic rationality as instrumental rationality: A critique. *Philosophy and Phenomenological Research*, 66(3), 612–640.
- Kolodny 2003. Love as valuing a relationship. *Philosophical Review*, 112(2): 135-189.
- Kornblith, H. 1993, Epistemic normativity. *Synthese*, 94, 357–376.
- _____ 2002 *Knowledge and its place in nature*. New York: OUP.
- Korsgaard, K. 1983. Two distinctions in goodness. *Philosophical Review* 92 (2):169-195.
- McHugh, C. 2012. Belief and aims. *Philosophical Studies*, 160(3), 425–439.
- Mercier, H. 2017. How gullible are we? A review of the evidence from psychology and social science. *Review of General Psychology*, 21: 103–122.
- Moore, G.E. 1903. *Principia Ethica*. Mineola: Dover Publications
- Levy, N. and Alfano, M. 2020. Knowledge from vice: deeply social epistemology. *Mind*, 129(515): 887–915.
- Palermos, S.O. 2016. Spreading the credit: virtue reliabilism and weak epistemic anti-individualism. *Erkenntnis*, 81: 305–334.
- Parfit, D. 2011. *On What Matters, Vol 1*. New York: OUP.
- Scanlon, T.M. 1998. *What We Owe to Each Other*. Cambridge, MA: Harvard.
- Scott, J. 2023. Should epistemic instrumentalists be more social? *Synthese* 201 (4):1-20.
- Steglich-Petersen, A. 2018. Epistemic instrumentalism, permissibility, and reasons for belief. In McHugh, Way & Whiting (eds.), *Normativity: Epistemic and Practical*. Oxford University Press.
- Schmidt, S. 2024. Epistemic blame and the normativity of evidence. *Erkenntnis*, 89 (1):1-24.
- Shah, N. 2003. How truth governs belief. *Philosophical Review*, 112(4), 447–482.
- Shah, N. and Velleman, D. 2005. Doxastic deliberation. *Philosophical Review* 114 (4):497-534.
- Sharadin, N. 2018. Epistemic instrumentalism and the reason to believe in accord with the evidence. *Synthese* 195 (9):3791-3809

- Simion, M. 2019. Knowledge-first functionalism. *Philosophical Issues*, 29 (1):254-267.
- Sperber, D. et al. 2010. Epistemic Vigilance. *Mind & Language*, 25: 359–393.
- Stocker, M. 1981. The limits of teleology and the ends of friendship. *The Journal of Philosophy*, 78(12): 747-765.
- Sylvan, K. 2020. An epistemic non consequentialism. *Philosophical Review*, 129 (1):1-51.
- Way, J. 2013. Value and reasons to favour. *Oxford Studies in Metaethics*, 8.
- Wedgwood, R. 2007. *The nature of normativity*. Oxford: Oxford University Press.
- Willoughby, J.B. 2022. Embedded epistemic instrumentalism: an account of epistemic normativity. *Philosophers' Imprint*, 22.
- Wei, X. Forthcoming. A practice-based account of the truth-norm of belief. *Episteme*.
- Worsnip, A. 2022. Epistemic normativity is independent of our goals. In Blake Roeber, Matthias Steup, Ernest Sosa & John Turri (eds.), *Contemporary Debates in Epistemology (3rd ed.)*. Wiley-Blackwell.