

ARTICLE

Epistemic Complicity

Cameron Boulton 

Department of Philosophy, Brandon University Faculty of Arts, Brandon, MB R7A 6A9, Canada
Email: boulton@brandonu.ca

(Received 22 January 2024; accepted 22 January 2024; first published online 29 February 2024)

Abstract

There is a widely accepted distinction between being directly responsible for a wrongdoing versus being somehow indirectly or vicariously responsible for the wrongdoing of another person or collective. Often this is couched in analyses of complicity, and complicity's role in the relationship between individual and collective wrongdoing. Complicity is important because, *inter alia*, it allows us to make sense of individuals who may be blameless or blameworthy to a relatively low degree for their immediate conduct, but are nevertheless blameworthy to a higher degree for their implication in some larger (or another person's) wrongdoing. In this paper, I argue that there is a distinctively epistemic kind of complicity. First, I motivate the distinction between direct and vicarious responsibility with three interlocking arguments, respectively appealing to: (i) the structure of rational agency; (ii) linguistic considerations; (iii) the role of 'principal' vs. 'accomplice' in legal doctrine. I show how these arguments naturally extend to the epistemic domain, motivating an epistemic form of vicarious responsibility. I then examine complicity as a mechanism of vicarious epistemic responsibility. To fill this out, I engage with an epistemic analogue of the debate about the role of intention versus causal contribution in complicity. I propose a Casual Account of Epistemic Complicity, arguing that it accommodates a wide range of cases, and enables fine-grained explanations of degrees of culpability for epistemic complicity. With an adequate account of epistemic complicity on hand, we can explain what is objectionable about an important class of epistemic agent who, on an individual level, may be epistemically blameless or blameworthy to a relatively low degree, but whose relation to other individuals or collectives nevertheless makes them epistemically blameworthy to a higher degree. I explore some broader implications of this result.

Keywords: collective epistemology; collective responsibility; complicity; vicarious responsibility; epistemic normativity; epistemic responsibility; ethics of belief; epistemic blame

1. Introduction

Consider two scenarios. In one, Jim forgets to secure the lid on his garbage bin and the raccoons make a big mess of his neighbour's yard. Other things being equal, it seems fair for the neighbour to hold Jim accountable for the mess. In the other scenario, Jim's 10-year-old son – whose job is to take care of the household garbage – forgets

© The Author(s), 2024. Published by Cambridge University Press. This is an Open Access article, distributed under the terms of the Creative Commons Attribution licence (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted re-use, distribution and reproduction, provided the original article is properly cited.

to secure the lid on the bin and the racoons make a big mess of the neighbour's yard. Even though the mess is more directly Jim's son's fault, it still seems fair for the neighbour to hold Jim accountable for the mess, even if not in quite the same way, or to the same degree.

These two scenarios highlight an important distinction. There seems to be a difference between being *directly* responsible for a wrongdoing, versus being *indirectly* or what I will call *vicariously* responsible for someone else's wrongdoing.¹ Parents can be vicariously responsible for things their children do. Something similar might be said of bosses and their employees. In a different but related way, an individual member of an army battalion might be vicariously responsible, at least to some degree, for atrocities caused by the *battalion*.^{2,3} A longstanding and still-growing literature in moral philosophy focuses on this distinction in a variety of guises.⁴ John Gardner goes so far as to argue that it marks two fundamentally different *parts* of morality:

I am responsible for what I do, and you are responsible for what you do. But on any credible view I need to give attention, in what I do, to what you will do in consequence. And you need to give attention, in what you do, to what I will do in consequence. In that sense, there are two parts of morality. There is what I should do simpliciter, and then there is what I should do by way of contribution to what you do (Gardner 2007: 9).

Regardless of whether Gardner is right that the distinction carves out two fundamentally different parts of morality, many agree there is something important in the neighbourhood. The minimal idea I am concerned with is, roughly, the idea that an agent can sometimes be responsible in a distinctive way for something that does not (solely) depend on their own conduct, but also on the conduct of another agent(s) (Glavaničová and Pascucci 2022: 1).

Much of the discussion – and the part that I will be most interested in – has taken place in the context of analyses of *complicity*, and complicity's role in the relationship between individual and collective wrongdoing. As authors like Christopher Kutz have argued, complicity is important because it allows us to make sense of individual agents who may not be particularly blameworthy for anything they have directly done, and yet seem blameworthy (perhaps to a high degree) for their implication in some larger (or

¹Differences in terminology can be found in the literature, with some authors – primarily in legal philosophy – reserving 'vicarious responsibility' for something altogether different from 'indirect responsibility' (Gardner 2007: 5; Kadish 1985: 338). I treat these terms synonymously and will primarily use 'vicarious'. As I explain below, I use this term, roughly, to refer to agents who are responsible for something that does not (solely) depend on their own conduct, but also on the conduct of another agent(s) (Glavaničová and Pascucci 2022: 1).

²See Goetze (2021: 211) for a helpful table categorizing *seven* different kinds of vicarious responsibility, including examples of each.

³So-called 'tracing cases' may illustrate a phenomenon closely related to vicarious responsibility. In tracing-cases, a time-slice of a single individual is in some sense responsible for the actions of a later time-slice of that same individual, even though the later time-slice lacks agential capacities or knowledge typically thought necessary for responsibility. In a sense, we might think of such cases as a kind of individual vicarious responsibility (cf. Blomberg 2023: 23).

⁴Some classic examples include Kadish 1985; Kutz 2000, 2007; Gardner 2007; Driver 2015; Lepora and Goodin 2013; Lawson 2013; Moore 2007. More recent work includes Blomberg 2023; Bennett 2021; Bazargan-Forward 2017; Björnsson 2020; Collins and De Haan 2021; Galvin and Harris 2023; Goetze 2021; Katz 2023; Mellor 2021; Glavaničová and Pascucci 2022; Zakaras 2018.

another person's) wrongdoing. In this paper, I want to argue that complicity can do something similar for us in the *epistemic* domain. It can do so by allowing us to explain what is objectionable about an important class of epistemic agent who, on an individual level, may be epistemically blameless or blameworthy to a relatively low degree, but whose relation to other individuals or collectives nevertheless makes them epistemically blameworthy to a higher degree.

Considering the predominance of work in epistemology examining parallels between the epistemic and moral domains, it is interesting that very little has been said about complicity, and the more general distinction between direct and vicarious responsibility, in an epistemic context. A great deal has been written about individual epistemic responsibility.⁵ And more recently, work has been done on *collective* and *shared* epistemic responsibility.⁶ But the notion of complicity, and vicarious responsibility more generally, cuts across both topics. As we'll see, epistemic complicity is an underexplored phenomenon often found precisely at the intersection of these domains.⁷

While some closely related issues are starting to receive attention, the overall trend has left important gaps. This paper aims to break new ground in a bid towards filling those gaps. Can people be vicariously responsible in the epistemic domain? What exactly does this amount to? I will argue people can be vicariously responsible in a distinctively epistemic way, and that epistemic complicity is a central mechanism of vicarious epistemic responsibility. I will defend this claim by motivating the idea that we should *expect* a kind of complicity in the epistemic domain, and developing an account of epistemic complicity, which I call the Causal Account.

The paper proceeds as follows. In Section 2, I defend the general distinction between direct and vicarious responsibility, relying on three mutually supporting arguments – one turns on claims about the structure of rational agency, one on linguistic considerations, and the other on the role of 'principal' and 'accomplice' in legal doctrine. As we'll see, motivating this general distinction is useful for developing a defence of the idea that there is a uniquely epistemic kind of complicity. In Section 3, I argue that considerations in support of the distinction between direct and vicarious responsibility readily extend to the epistemic domain, motivating an examination of 'vicarious epistemic responsibility'. In Section 4, I examine one way people might be vicariously epistemically responsible, namely by being *complicit* in another agent or collective's epistemic failing. To fill this out, I engage with an epistemic analogue of the debate about the role of *causal contribution* versus *intentions* or *shared goals* as minimal conditions on complicity (Section 4.1). Drawing on work in this debate, I propose an account of epistemic complicity I call the Causal Account (Section 4.2); I highlight some of its advantages, including its ability to provide a fine-grained explanation of degrees of culpability for complicity. In section 5, I defend the account against the problem of 'causeless complicity'. My primary concern in all this is to motivate a compelling option for thinking about epistemic complicity, and deploy it in two main ways: (i) to explain what it takes for an account of complicity to be *epistemic*; and (ii) to demonstrate the theoretical usefulness of epistemic complicity. To this latter end, in Section 6, I argue that

⁵For some central examples, see Cassam 2019; Levy 2007; Nottelmann 2007; Montmarquet 1992; Owens 2000; Peels 2016; Van Woudenberg 2009; Schmidt and Ernst 2020; Zagzebski 1996. There are too many others to list here.

⁶See Schwenkenbecher 2022; Fleisher and Šešelja 2023; De Haan 2021; Mitova 2022; Palermos 2022.

⁷Two notable recent exceptions of authors engaging with vicarious epistemic responsibility – though not in this terminology, and in more specific ways than I do so here – are Mitova (2022) and Palermos (2022).

epistemic complicity has an important role to play alongside ‘shared epistemic responsibility’ (Millar 2021) in our epistemic blaming practices.

2. Vicarious responsibility

The term ‘responsibility’ is multiply ambiguous. Philosophers sometimes distinguish between responsibility as *attributability* (Shoemaker 2011), *answerability* (Smith 2008, 2012) and *accountability* (Scanlon 1998, 2008; Smith 2008). We can also distinguish between *forward-looking* and *backward-looking* uses of ‘responsible’, where the former is used, roughly, to denote that someone is *owing* something, and the latter is used to denote that someone has *done* something. In this paper, I use ‘responsible’ in a backward-looking way, and in the sense of ‘accountable’, where I understand accountability to imply openness to blame. So, when I say X is responsible for Y, I mean to say that X has *done* something (in a broad sense of ‘do’) for which they may be an appropriate target of blame.

Not everyone agrees that there is such thing as ‘vicarious responsibility’. H.D. Lewis writes: ‘If I were asked to put forward an ethical principle which I considered to be especially certain, it would be that no one can be responsible, in the properly ethical sense, for the conduct of another. Responsibility belongs essentially to the individual’ (Lewis 1948). Appearing to agree, Joel Feinberg (1966) says, ‘*there can be no such thing as vicarious guilt*’ (676, original italics). Regardless of whether we find these bold statements compelling, it does seem fair to question whether we need to draw a fundamental distinction between two kinds or ways of being responsible along these lines. After all, even if I am partly responsible for *your* X-ing, it’s nevertheless true that I’m directly responsible for *something* – namely, whatever contribution I made to your X-ing. Perhaps the story of my responsibility ends there.⁸

In this section, I consider three arguments in support of the general distinction between direct and vicarious responsibility. These are (i) an argument from the structure of rational agency; (ii) an argument from linguistic considerations; and (iii) an argument based on the distinction between ‘principal’ and ‘accomplice’ in legal doctrine.

2.1. The structure of rational agency

According to John Gardner (2007), the distinction between direct and vicarious responsibility⁹ is baked into the structure of rational agency. Consider two ways normative reasons might be said to be *personal*, or to have normative force specifically for one agent rather than another. Reasons may be personal ‘in respect of conformity’, or personal ‘in respect of attention’ (Gardner 2007: 8). When I make a promise to you, this gives me a pro tanto reason to keep my promise. In an important sense, only *I* can flout this reason. But that doesn’t mean that other people are incapable of *appreciating* my reason to keep the promise. Perhaps under the right circumstances, any number of people owe my reason the same rational attention that I do. As Gardner puts it: ‘There are many possible ways for other people to give my reason the full attention it deserves, assuming that the reason is impersonal in respect of attention. The only principled limit to what they can do to help me is that they can’t actually keep the promise for me’ (8).

⁸See also Bazargan-Forward (2017) and Moore (2007) for critical discussions of the distinction, or at least closely related ideas.

⁹His term is ‘indirect’ responsibility.

My reason to keep my promise is personal in respect of conformity, but that doesn't preclude it from being *impersonal* in respect of attention. These observations about the structure of rational agency can be used to motivate the distinction between direct and vicarious responsibility. If a person is responsible for complying with those reasons which are personal in respect of conformity, it seems they are responsible for so complying in a way that no one else can be. This just goes hand in hand with what it means to say that these reasons are 'personal' (only *you* can break your promise). So, if an agent bears some culpability for the failure of some *other* agent to comply with (their own) reasons which are personal in respect of conformity, it seems they must be so responsible in a different kind of way. But since we clearly can bear some culpability for such failures of others (consider the examples from the Introduction), it seems there is a different way of being responsible than the first. This is the distinction between direct and vicarious responsibility.

Gardner is concerned to show that adopting a distinction between direct and vicarious responsibility does not commit one to 'agent-relativism' about reasons, or the idea, roughly, that there are some normative reasons which fundamentally bear on what *I* should do as opposed to anyone else (Hammerton 2018; Nagel 1986; Parfit 1984; Pettit 1987). The distinction between agent-neutralism and agent-relativism is widely regarded as one of the most important debates in ethics.¹⁰ It is often considered an essential point of contention – a way *framing* the disagreement – between deontologists and consequentialists. It would be unfortunate if only those committed to one side of this debate – arguably the less dominant side (agent-relativism) – are rationally able to endorse the distinction between direct and vicarious responsibility. But on closer inspection it is clear that the distinction between direct and vicarious responsibility cuts across this debate. The debate is not about whether there are at least some reasons which are personal in respect of conformity – this is simply a conceptual truth about reasons that all parties to the debate accept. It is about whether there are at least some reasons which are personal in respect of *attention*. Roughly speaking, the relativist says 'yes', and the neutralist says 'no'. Thus, regardless of whether one is a consequentialist, deontologist, agent-neutralist or agent-relativist: 'the distinction between principals and accomplices is embedded in the structure of rational agency. As rational beings we cannot live without it' (2007: 9).

2.2. Linguistic considerations

Consider eating, having sex or going for a run. Sandy Kadish (1985) calls such actions 'non-proxyable'. Nobody else can eat, have sex or go for a run 'through you', or in other words by making a causal contribution to your actions. Notice that certain morally laden terms also appear, at least on very natural readings, to refer to actions that are non-proxyable: raping, torturing, killing, betraying and many others (Kadish 1985). These too seem like things that nobody else can, strictly speaking, *do* 'through you', i.e. by making a causal contribution to your actions.¹¹ Our moral vocabulary encodes

¹⁰See Hurka (2003: 628), Ridge (2005/2022) and Portmore (2013).

¹¹For those with differing linguistic intuitions about some of these verbs, such as 'killing', consider how natural it can sound to say that someone was the 'killer, in effect'. When that person intentionally caused a death through the causal contribution of someone else, it can sound natural to call them the 'killer, in effect'. To my mind, it is more natural to call the one who is directly causally connected to the death the *killer*, and the one who brought about the death through such an intermediary the 'killer, in effect'. As Gardner points out, a killer is not the same thing as a killer in effect (2007: 9).

a division in the causal profile of certain wrongs. As such, it appears to reveal a deep concern for the difference between an individual's *direct* connection to a wrongdoing, versus (or in addition to) an individual's connection to a wrongdoing through the actions of someone else or others. In other words, it appears to reveal a concern for the distinction between direct and vicarious responsibility.

Consider the distinction between killing and intentionally causing someone's death (say, by hiring a hitman), or raping and intentionally bringing about non-consensual sex (say by forcing someone to rape another). For any number of related items of vocabulary, a structurally similar difference in causal profile underpins these distinctions, seeming to correspond nicely with the distinction between direct and vicarious responsibility. One can be responsible for raping someone, but one can also be responsible for intentionally bringing about non-consensual sex. In the latter case, in addition to being directly responsible for intentionally bringing about non-consensual sex, one is also indirectly or vicariously responsible for rape – after all, that is, at least in part, what non-consensual sex amounts to. We tend to reserve 'rapist' for the person who is *directly* responsible for raping. Gardner frames this in terms of a difference in how 'causally refined' a given action-type is. Rape is a causally refined way of intentionally causing non-consensual sex; intentionally causing non-consensual sex is a causally refined way of contributing to a rape. Our moral vocabulary marks a distinction between ways agents can be involved in, or causally contribute to, a wrongdoing: directly or indirectly/vicariously.¹² Of course, our moral categories could be fundamentally confused. But these observations suggest the distinction between direct and vicarious responsibility is central to our moral lives, and so a legitimate one to embrace for theoretical purposes.¹³

2.3. 'Principal' vs. 'accomplice' in legal doctrine

As many legal theorists have argued, we need a distinction between principal and accomplice – sometimes referred to as the 'doctrine of complicity' – in order to coherently and justifiably make decisions about the wide-ranging cases we face in law (Kadish 1985). Consider international tribunals where high-ranking military commanders are tried for their role in mass killings of innocents. Or consider a case in which Tony Soprano is being tried for ordering his lackeys to rob a bank. Even though neither the commanders nor Soprano personally committed the crimes (the actual killings, the actual robbing), they clearly played an important – essential – role. We need a way of fully holding such actors legally to account.¹⁴ Without a distinction between the kind of liability at issue when one personally or directly commits a crime, and a kind of liability at issue when one is somehow involved despite not directly committing the crime, it is difficult to see how we could do so. The principal-accomplice distinction is deployed for these purposes.

¹²There is room for debate about whether vicarious responsibility can ultimately be understood in terms of *causal* contribution at all. Gardner's view on this is controversial. We will return to this issue in Section 4. What matters for present purposes is merely the observation that there appears to be a fundamental difference in how agents can be *connected* (whether causally or otherwise) to wrongdoings, marked out in our moral vocabulary.

¹³See Bazargan-Forward (2017: 7) for an opposing view on the relevance of non-proxyable actions for the distinction between direct and vicarious liability in a legal context.

¹⁴Of course, these actors may have perpetrated their own crimes through their involvement in the crimes of others. But in such cases, this often does not exhaust the extent of their liability.

The cases above are interesting in that, in both, we might think the ones who did not directly commit the crimes are in fact *more* liable than the ones who did. Some theories of accomplice law are designed to treat such actors as *principals* rather than accomplices, for this very reason.¹⁵ While ubiquitous, the ‘doctrine of complicity’ and issues around how to understand accomplice liability are controversial and notoriously complex. However, the complexity concerns how to consistently justify rulings across various types of cases (sometimes highly complex ones). The issue is not *whether* the concept of accomplice liability is useful and important – seemingly indispensable – in legal theory. Thus, the point remains: without some apparatus for consistently assigning liability to those clearly in some way involved in crimes, despite not perpetrating the crimes themselves, we would be unable to justifiably and fully hold such bad actors to account.¹⁶ The distinction between principal and accomplice liability seems to presuppose a distinction very much like our distinction between direct and indirect or vicarious responsibility. It seems designed to help us deal legally with these different modes of responsibility. Thus, it appears that our distinction is presupposed by an indispensable element of legal doctrine. In my view, this is a significant reason to take the distinction seriously.

The indispensability of the principal-accomplice distinction in legal doctrine might seem like a fairly defeasible argument in support of the distinction between direct and vicarious responsibility. Legal responsibility is not the same thing as moral responsibility. I have three things to say in response to this. First, the law is a central facet of our social, political and public lives – one arguably characterized, in part, by its role in codifying and institutionalizing our responsibility practices. We can always evaluate the law from other normative points of view. But other things being equal, if a mechanism designed to codify and institutionalize our responsibility practices marks a fundamental distinction between two kinds of liability (responsibility), it seems fair to expect something like this distinction to be an important part of responsibility more generally. Second, as Chiara Lepora and Robert Goodin point out in their own defence of appeals to legal doctrine, ‘[legal doctrine] is where we find the most *thorough* discussions of complicity to date, and when moral philosophers discuss these matters they most naturally turn to those discussions to see what they might usefully borrow’ (Lepora and Goodin 2013: 89 italics mine). Third, there is a striking degree of corroboration between the argument from legal doctrine, the argument from linguistic considerations and the argument from rational agency. The argument from rational agency might even contain the beginnings of an *explanation* of the observations we’ve made in these last two subsections. While each argument is of course defeasible, taken together they comprise considerable support for the distinction between direct and vicarious responsibility.

¹⁵Some of the major theories of accomplice law include common law, which treats principals as the ones who commit the act (everyone else is mere accomplice); widely applied in German courts for decades was the view that principals are those with the culpable *mental* state, the intent to carry out the crime. There are also numerous views in between, such as the Control Theory developed by German criminal law theorist Claus Roxin.

¹⁶See Moore (2007) for in-depth critical discussion and an apparently dissenting view. Moore himself notes that it is common in Anglo-American law to distinguish between *vicarious* liability and *accomplice* liability as distinct kinds of liability – accomplice liability being a broader or less stringent notion. Moore does not challenge the more stringent notion of vicarious liability, and so, he would seem to agree with the basic spirit of my argument in this section that, at a minimum, some kind of distinction between direct and vicarious liability (responsibility) is an essential part of legal doctrine.

3. Vicarious responsibility in the epistemic domain

My primary concern is to examine possible mechanisms of vicarious responsibility in the *epistemic* domain. If the distinction between direct and vicarious responsibility is so fundamental in the moral domain, might we also find it in the epistemic domain? And if so, what are its mechanisms?

Before getting to this, just as the term ‘responsibility’ is multiply ambiguous, it’s important to be clear on some different things we might mean by ‘responsibility in the epistemic domain’, or even ‘epistemic responsibility’, whether vicarious or otherwise (and keeping in mind that I am using ‘responsibility’ in a backward-looking, accountability sense). Crucially, these phrases are ambiguous between two domains of responsibility. On one hand, ‘responsibility in the epistemic domain’ might denote a distinctively epistemic kind of responsibility – something uniquely oriented around respect for epistemic norms and values, and thus (perhaps) importantly different from moral responsibility. On the other hand, ‘responsibility in the epistemic domain’ might simply denote a species of moral responsibility, one that contingently takes epistemic failings as its object.

One way to delineate between these two domains of responsibility is as follows:

- An agent or collective X is *epistemically* responsible for an epistemic failing Y just in case, barring excuse, X is open to epistemic blame for Y.
- An agent or collective X is (morally, practically, etc.) responsible for an epistemic failing Y just in case, barring excuse, X is open to (moral, practical, etc.) blame for something brought about, or partially constituted by Y.

Assume there is a generic type of response we call ‘blame’ which can sometimes be directed towards moral failings, sometimes epistemic ones and sometimes others. When we speak of ‘types’ of blame, this can be defined in terms of the most salient *failing* a blame response concerns (cf. Boulton 2023, 2021; Brown 2020a, 2020b; Piovarchy 2021; Schmidt 2021). When I blame my boss for holding prejudiced attitudes against me, if a moral failing is most saliently at issue, then this is moral blame – despite the fact that his prejudice is partially constituted by an epistemic failing. When I blame my brother for hastily concluding he’s about to win the lottery, if an epistemic failing is most saliently at issue, then this is epistemic blame. We can thus differentiate our senses of ‘responsibility in the epistemic domain’ in terms of whether appropriate blame is most saliently directed at the epistemic as opposed to moral dimensions of whatever failing is at issue.

The two domains of responsibility are not mutually exclusive. Someone can be both epistemically and morally responsible for the same epistemic failing. Indeed, this is most often the case. But that doesn’t mean these two domains of responsibility do not mark out importantly distinct dimensions of our interpersonal and normative lives. All of this is crucial because it seems much more interesting to examine whether, and in what ways, people can be vicariously *epistemically* responsible in the epistemic domain, where what we have in mind is a distinctively epistemic kind of responsibility. Focusing merely on vicarious (moral) responsibility for epistemic failings, our discussion may simply end up being a kind of redundant instance of the general discussion about vicarious responsibility, a discussion that has been going on for years in the moral domain.

So, can we be vicariously epistemically responsible? The first thing to note is that there seems to be no principled barrier to extending the arguments for vicarious moral responsibility considered in Section 2 to the epistemic domain. Consider the argument from the structure of rational agency. That argument turns on claims about different ways that

reasons can be *personal*. There is no less plausibly a distinction between *personal reasons in respect of conformity* and *impersonal reasons in respect of attention* in the epistemic domain. That I am attending to a red apple in good lighting may give me an epistemic reason to believe I am looking at an apple. It seems true that only *I* am capable of flouting this reason, even if it is also true that indefinitely many other people are perfectly capable of appreciating the force of this reason, and perhaps even have their own reasons to get me to appreciate it. So, it seems (only) I have a personal epistemic reason in respect of conformity here, where others have (at most) an impersonal epistemic reason in respect of attention.

If a person is epistemically responsible for complying with those epistemic reasons which are personal in respect of conformity, it seems they are epistemically responsible for so complying in a way that no one else can be. So, if an agent bears some culpability for the failure of some *other* agent to comply with (their own) epistemic reasons which are personal in respect of conformity, it seems they must be epistemically responsible in a different kind of way. So, if we can bear some culpability for such failures of others, it seems to follow there is a different way of being epistemically responsible than the first. This is the distinction between direct and vicarious epistemic responsibility.¹⁷

Of course, a crucial question here is whether people can bear some culpability for the failures of others in the epistemic domain. And more to the point: can people bear

¹⁷Might the transparency of belief disrupt the analogy between epistemic reasons that are personal in respect of conformity and practical reasons that are personal in respect of conformity? Suppose I become aware that you have an epistemic reason to believe that *p*. Don't I thereby gain an epistemic reason to believe that *p*? If so, perhaps there is an important difference between the normative domains in this regard, and perhaps this poses a challenge for extending the argument from the structure of rational agency to the epistemic domain. In response, I do not think this observation poses a challenge. First, there are cases in which I become aware that you have an epistemic reason to believe that *p* where I nevertheless do not myself thereby obtain an epistemic reason to believe that *p*. Perhaps I'm watching the Truman Show, and I can see that Truman has all sorts of epistemic reasons to believe things about the world around him, while nevertheless appropriately refraining from believing those things myself. Second, even if there are many cases in which I gain my own epistemic reasons by learning about the epistemic reasons you have, this does not undermine the idea that only you are capable of flouting your epistemic reasons which are personal in respect of conformity, and only I am capable of flouting mine. As long as there are also cases where an agent bears some culpability for the failure of some *other* agent to comply with (their own) epistemic reasons which are personal in respect of conformity, this is enough to get the argument from the structure of rational (epistemic) agency off the ground. Thanks to Sarah Wright for pressing this worry.

some *epistemic* culpability for such failures of others in the epistemic domain? Before getting to these questions, it's worth noting that our other two arguments from Section 2 seem equally extendable to the epistemic domain. Just as moral vocabulary encodes causal differences in our very conception of certain wrongdoings, so too does epistemic vocabulary. Consider: no one else can form an unjustified belief *through me*, by making a causal contribution to my actions or mental states. Regarding the principal-accomplice distinction in legal doctrine, this central feature of legal doctrine seems equally defeasibly relevant to our responsibility practices in both the moral and epistemic domains.

So, again, there appears to be no principled prohibition on a distinction between direct and vicarious responsibility in the epistemic domain, even understood as a distinctively epistemic kind of responsibility. What remains is to establish that people can bear epistemic culpability for the direct epistemic failings of others, and to articulate the mechanisms by which this takes place.

4. Complicity as a mechanism of vicarious responsibility in the epistemic domain

I want to argue that one central mechanism of vicarious epistemic responsibility is epistemic *complicity*. What exactly is complicity? And how can it make one person vicariously responsible for the wrongdoing of another?

Two initial structural points that will be important below are: (i) complicity is not (only) a wrongdoing in its own right, it is a *way of participating or being involved* in another's wrongdoing; and (ii) complicity does not preclude principal wrongdoings on behalf of the complicitor in their involvement in another's wrongdoing. I may commit any number of my own wrongdoings through my involvement in your wrongdoings, though I need not do so. As a way of being involved in another's wrongdoing, 'complicity' comes in many varieties and has a wide range of conceptual cousins. Full-joint wrongdoing, co-operation, collaboration, conspiracy, condoning and conniving have all been regarded as forms of complicity, or closely related phenomena. And there are of course important differences between all of these things. One interesting question is whether an equally rich array of *types* of complicity can exist in the epistemic domain. But rather than addressing that issue here, my aim going forward will be to examine some candidate *minimal conditions* on complicity in the epistemic domain. Locating a minimal condition on complicity in the epistemic domain would explain what various types of complicity have in common, and would thus go a long way towards gaining a better understanding of complicity in the epistemic domain.

4.1. The minimal condition debate

Approaches to moral complicity tend to emphasize one of two criteria as minimal conditions on complicity. *Causal* views require a kind of causal contribution to a principal wrongdoing as a minimal condition on complicity. *Intentional* views require a certain kind of *attitude, intention* or even *shared goal* vis-à-vis the principal wrongdoer as a minimal condition on complicity.

The debate between these views tends to proceed by one side or the other emphasizing the importance of their preferred condition(s), and then arguing that their opponent's condition(s) are not necessary for complicity. For example, Christopher Kutz (2000, 2007), the most well-known proponent of what I'm calling the intentional view, argues that complicity's minimal condition is a kind of 'participatory intention'. This is an intention to do one's part of a collective act, where one's part is the task one should perform if the collective is to be successful in realizing a shared goal (Kutz 2000: 81). On this view, what matters for complicity is that a complicitor has some kind of shared conception with a principal wrongdoer of *what they are doing*, and *aims to do what they can* to contribute *their part* to the principal wrongdoing.¹⁸

¹⁸Kutz's approach starts from a minimal account of *joint action* – minimal in the sense that it avoids robust conditions on joint action, such as those developed by Michael Bratman (1993) and Margaret Gilbert (2000, 2002): conditions involving common knowledge and wholly shared conceptions of the joint act. Kutz instead focuses on bare criteria needed to capture groups of people who are intuitively collectively responsible for certain harms. As Brian Lawson (2013: 232) puts it, the account of joint action is *normative*: it starts from intuitive claims about who should be held accountable, and works from there towards minimal criteria on what unifies those accountable. According to those minimal criteria, a set of individuals is engaged in a joint action when each individual (i) has a sufficiently overlapping conception of a shared goal, G, and (ii) each performs a voluntary action in the aim of contributing their part toward realizing G (Kutz 2000: 90). Kutz's account of complicity drops out of his minimal account of joint action.

Kutz argues that an intentional approach is crucial because of the way complicity can sometimes be entirely ‘causeless’, or in other words, because of the importance of making sense of individuals who seem not to make any meaningful difference to a greater principal wrongdoing, and yet somehow seem implicated in an important way. Kutz’s central example is the WWII bombing of Dresden. The crews who bombed the city had a sufficiently overlapping conception of a shared goal – carpet bombing the city – and each did their part in contributing to the achievement of that goal: loaded their bombs, flew their planes to the location and dropped their bombs. Each individual crew member is complicit in the firestorm that destroyed the city, not because they made some individual difference to the firestorm,¹⁹ but because of the ‘structure of their wills’ (Kutz 2000: 141).

Approaching complicity from the other direction, the most prominent defenders of a causal account, Chiara Lepora and Robert Goodin (2013), have argued at length that shared goals or ‘participatory intentions’ are likewise inessential for complicity (cf. Gardner 2007; Petersson 2013). Instead, despite Kutz’s worries about causeless complicity, according to Lepora and Goodin making a certain kind of causal contribution to a principal wrongdoing is indeed a minimal condition on complicity. There are of course a wide range of competing views about the role for causation and participatory intentions in our understanding of complicity (Bazargan-Forward 2017; Blomberg and Hindriks 2020; French 2016; Gardner 2007; Katz 2023; Lepora and Goodin 2013, 2016; Petersson 2013). Given the complexity of this literature, I cannot hope to conclusively defend one sort of approach over another, nor do I regard doing so as necessary for my bigger-picture project of identifying mechanisms of vicarious responsibility in the epistemic domain. However, a causal approach strikes me an especially helpful starting point for such a project, in no small part because I agree that participatory intentions or shared goals seem absent in many intuitive cases of complicity in both the moral and epistemic domains (as I’ll discuss in Section 4.2). I also think there are significant resources for dealing with cases of ‘causeless complicity’ (I’ll discuss this in Section 5). I will ultimately remain open on the question of whether there are kinds of complicity in the epistemic domain which do not involve any kind of causal contribution. But my strategy going forward will be to argue there is sufficient motivation for the causal framework to give it an important role in our understanding of epistemic complicity, and to highlight what it can do for us in the epistemic domain.

4.2. The causal account of epistemic complicity

Consider a man who has devised a clever plan to keep his husband asleep with sleeping pills while he gallivants around town with a secret lover. The pills are prescription, so he

Sometimes Kutz explicitly emphasizes a role for even more deflationary kinds of ‘participatory intentions’. Indeed, one of his aims in approaching complicity through a *minimal* account of joint action is to be able to account for ‘intentional participation by cognitively vague, alienated, or dyspeptic agents [...] [to] make sense of collective action in familiar circumstances of routinized cooperation, hierarchical authority, and compartmentalized information’ (2000: 102). One source of pressure on Kutz’s account here, however, is that the more deflationary he goes on participatory intentions, the more he runs the risk of counting *too many* agents as complicitors. See Kutz (2000: Ch. 3.7). Thanks to Olle Blomberg for helpful discussion.

¹⁹Kutz regards the Dresden case as a paradigm example of individual participation in a wrong that is massively overdetermined: ‘The city was bombed in three raids, and at least 1000 planes and 8000 crewmen were directly involved in the raids, in various roles as pilots, navigators, bombers, and gunners. The firestorm was already raging before many crews dropped their bombs. Each crewman’s contribution to the conflagration, indeed, each plane’s, was marginal to the point of insignificance’ (2000: 118).

needs to get them from his doctor (Lepora and Goodin 2013: 81–83). The doctor might have a variety of intentions in the process of their interactions around the pills. In one sort of case, the doctor is himself the lover. Here, the two have sufficiently overlapping *participatory* intentions and count as complicit in the full Kutzian sense.

In another case, the doctor is conducting a study on these new sleeping pills and has ways of gathering data through the man's shady use of the pills. He knows what the man is up to, and doesn't approve, but it's such a valuable source of data he prescribes the pills anyway. Here, there are overlapping intentions: the man's achieving his aim requires an act (putting his husband to sleep) that is merely a means to the doctor's independent aims (learning more about the pills).

Following Lepora and Goodin, we might say that this is an overlapping of *contributory* intentions, rather than *participatory* ones (2013, 82). So even here we fail to satisfy the conditions of Kutzian complicity, despite a kind of interconnectedness between the 'wills' of these two people. And yet, it seems clear that the doctor is complicit in the man's shady activities.

We can go further still and imagine a case where the doctor knows what the man is up to, does not approve, but prescribes the pills anyway, simply because he wants to make a buck. There is clearly no shared goal in such a case. And yet it remains clear the doctor is complicit. In a final sort of case, the doctor does *not* know what the man is up to, nor can we reasonably expect him to have known. He prescribes the pills thinking they are for the man himself. In this case, importantly, the doctor does not seem complicit in the man's shady activities. And this provides us with an important epistemic constraint on a causal theory of complicity.

Across this range of cases, what seems to matter for our judgements of the doctor's complicity is not that he willingly participates in the man's aims, nor merely that he contributes to their success, but that he knowingly or at least culpably ignorantly contributes to them, while knowing or being culpably ignorant that they are wrong. Putting all this together, and following Lepora and Goodin, I propose the following:

Causal Account of Complicity: A secondary agent is complicit in a principal harm or wrongdoing only if: (i) they causally contribute to the harm or wrongdoing voluntarily or non-accidentally; (ii) they make the contribution knowingly, or culpably ignorantly; and (iii) they know or are culpably ignorant that the principal harm or wrongdoing they contribute to is wrong or harmful.²⁰

Note, importantly, this is a minimal condition on complicity. It is not intended to rule out that certain forms of complicity may involve more. To quote Lepora and Goodin: 'joining in the other's wrongful action or sharing in his wrongful purpose or intention are clearly even worse. But that would be much more than is minimally required for the situation to constitute one of complicity' (83).

Of course, my main goal is to examine how these ideas interact with the epistemic domain. Here is one way of extending this framework to the epistemic domain:

Causal Account of Epistemic Complicity: A secondary agent is epistemically complicit in a principal epistemic harm or wrongdoing only if: (i) they causally contribute to the epistemic harm or wrongdoing voluntarily or non-accidentally; (ii) they make the contribution knowingly, or culpably ignorantly; and (iii) they

²⁰This principle is based on Lepora and Goodin's preferred account (2013: 83).

know or are culpably ignorant that the principal epistemic harm or wrongdoing they contribute to is epistemically wrong or harmful.

The basic idea, of course, is to replace the kinds of wrongs or harms at issue in the original account (moral ones) with epistemic wrongs or harms. These can range from the formation of culpably unjustified beliefs, culpably faulty inquiry or perhaps certain kinds of epistemic or hermeneutical injustices (Fricker 2007; Hänel 2020; Jackson 2019; Medina 2012a, 2012b; Pohlhaus 2012). Is this a plausible way of thinking about ‘epistemic complicity’? One immediate point to make is that this account deals nicely with some interesting and important kinds of cases. I’ll begin with one for now, and return to some others later on.

Newscaster: Carl is a newscaster who goes on primetime every weeknight and constructs carefully crafted monologues designed to divide the nation. He knows he is contributing in significant ways to the degradation of the intellectual lives of a great many citizens. He’s contributing to their propensity to engage in epistemically blameworthy conduct. But he doesn’t really care. The more divided the nation is, the more ratings he gets. The more ratings he gets, the higher he gets paid. His sole motivation in constructing the monologues is profit.

Carl seems in some sense complicit, at least to some degree, in his audience’s epistemically blameworthy conduct. Note that this isn’t necessarily because he shares some kind of goal with his audience (in the ‘participatory’ sense). To be sure, Carl carefully designs his monologues to appear as though they have the aim of reaching the truth about politics in his country. And we might say that at least some of his audience tunes in with precisely that aim themselves – namely, of reaching the truth about politics in his country. So, it might appear as though Carl and his audience have some kind of shared epistemic aim of generating knowledge and true beliefs about politics in his country. However, the appearance of Carl’s participation in this aim is just a carefully crafted ruse. It is a *means* to generating ratings and securing higher profits. Whatever beliefs he manages to instil in his audience are not an end in themselves. Indeed, Carl actually disbelieves much of what he says, and can clearly see the flaws in the lines of reasoning he knows will seduce his audience. His only goal is profit, which happens to overlap with the goals of his audience only in the sense that their goals partly constitute a means to his ends.

It seems, then, that a Kutzian framework is incapable of explaining how Carl might be complicit in the intellectual failings of his audience. Meanwhile, the opposite seems true of the Causal Account. Carl makes a *contribution* to the collective epistemic harm constituted by the tendency to engage in culpable epistemic failings amongst his viewers, he knows that he so contributes, and he knows the principal epistemic harm he contributes to is epistemically harmful. According to the Causal Account, Carl is complicit in his audience’s culpable epistemic failings. And so, to this extent, he is vicariously *responsible* for their culpable epistemic failings. He shares in at least part of the blame (for *their* epistemic failings).

Is complicity really the right notion for Carl? Someone might object that he is simply directly responsible for doing objectionable things like lying and deceiving. Perhaps whatever involvement he has in the intellectual conduct of his *audience* is irrelevant in analysing what’s objectionable about Carl’s conduct. I disagree. Recall from earlier, it is of course possible for a complicitor to be the principal wrongdoer of all sorts of

harms or wrongs in their role as complicitor. From this platitude alone it does not follow that Carl is not also the appropriate target of further blame for his role in the failings of others. And that is what seems plausible here: in addition to being a liar and a deceiver, Carl is blameworthy for his role in the culpable epistemic failings of his audience.

A crucial feature of the case is that Carl does not believe anything he says. He just knows that the more he divides the nation, the higher his ratings will be. It might seem clear that there is something morally objectionable about Carl's activities (and so, equally clear that he's morally blameworthy). Is he also *epistemically* blameworthy? How so, exactly? This is a crucial question, because my *test* for whether we are talking about a distinctively epistemic kind of complicity is whether it is a form of *epistemic* responsibility, and the test for the latter is the appropriateness of epistemic blame. Does Carl pass that test? This might seem puzzling because, in not believing anything he says, nor actually endorsing many of the chains of reasoning he deploys, Carl may *himself* manage to avoid violating many plausible epistemic norms in constructing his monologues. Indeed, putting them together takes a lot of careful reasoning.²¹

But recall that moral and epistemic blame are not mutually exclusive. So, while it's clear that, other things being equal, Carl is open to moral blame for his approach to newscasting, it is also plausible that he can be appropriately blamed, at least to some degree, for the culpable epistemic failings he contributes to. There is a sense in which he is blameworthy that is not just moral in nature. Carl seems *in some sense* complicit in the culpable epistemic failings of his audience. But since complicity isn't (just) a wrong in its own right – it is a *way of being involved in* another's wrong – and since, in our case, the relevant principal culpable failings are epistemic, Carl is himself (vicariously) involved in culpable epistemic failings. Moreover, he seems *blameworthy* for that. Thus, on our approach, he counts as epistemically blameworthy. So, we have resources for explaining how Carl is vicariously *epistemically* responsible for a collective's culpable epistemic failings.

Crucially, then, the Causal Account of Epistemic Complicity does not merely help us to explain a trivial instance of the sort of thing moral philosophers have already done a great deal to illuminate in debates about complicity. It locates a genuinely novel phenomenon, something that is not identical with, nor reducible to, moral complicity nor vicarious responsibility simpliciter. This in turn is interesting, in my view, because it tells us something fundamental about the structure of the epistemic domain. It tells us that the epistemic domain itself admits of one person being epistemically responsible in a distinctive way for the culpable epistemic failing of another individual or collective. More than that, the Causal Account goes a significant way towards explaining how this is possible. I should emphasize that the idea I am interested in here is not that agents are either wholly blameless or blameworthy for their individual conduct, and that complicity changes this completely. Rather, my point is that there is an important

²¹Of course, on any plausible account of the norm of assertion, Carl violates that norm when speaking through his monologues. I'll return to the relationship between Carl's direct epistemic blameworthiness and his vicarious epistemic blameworthiness below. To avoid issues about the norm of assertion, it may be helpful to focus, in the first instance, on claims about Carl's *construction* of his monologues, and not the actual broadcasting of them. Alternatively, we might even focus on a variant of the case involving Carl's producer. Both of these suggestions may of course introduce other complications, so I will leave this to one side for present purposes.

role for complicity in understanding the full scope of a given individual's blameworthiness, both in the moral *and* epistemic domains.

An additional attractive feature of causal accounts of complicity is that they can provide fine-grained explanations of degrees of culpability for complicity. This is because causal contributions admit of fine-grained properties that can be used to specify differences between types and degrees of complicity. More than that, recall that causal accounts do not rule out a role for intentions in complicity. We can readily allow that cases involving shared goals and participatory intentions may very well be *worse* examples of complicity than cases that do not. In other words, whatever story a Kutzian framework offers about degrees of culpability may simply be imported as one dimension among others within the Causal Account.

Lepora and Goodin have developed a detailed framework for assessing degrees of culpability for complicity, one that divides considerations at the broadest level between: (i) facts about the causal role of the secondary agent, and (ii) facts about the agent's mental state. We can distinguish between a variety of properties within each of these dimensions, such as how *essential*, *central*, *proximal* and *irreversible* the agent's causal contribution was, as well as whether the agent had shared goals with the principal, and the degree of their authorship in any planning involved, and responsiveness to changes in others' intentions and actions throughout the execution of a principal wrongdoing (Lepora and Goodin 2013: Ch. 4). Moreover, given that complicity is not (just) a wrong in its own right, but rather a *way of being involved* in the wrong of another, we must also factor in considerations about how bad the relevant principal harm or wrongdoing is. If it is something really bad, then, even if the agent's causal role is relatively minimal (inessential, or non-central, for example) they may still be highly culpable for their complicity. Conversely, if the agent's causal role is highly central or essential, or if they have some kind of shared goal with the principal wrongdoer, then even if the principal wrongdoing is relatively small, so to speak, the complicitor may nevertheless be highly culpable. The overall degree of culpability for complicity will be a function of all these things combined.²²

5. Causeless complicity

All of this sounds good as far as it goes. But it leaves us with an important question. What do causal accounts have to say about the supposedly central and important phenomenon of 'causeless complicity'? This question seems no less pressing in the epistemic domain. Consider:

Disinformation campaign: Sometime during the height of the pandemic, a team of 10,000 anti-vaxxers assembles. The team is largely unified in its opposition to the use of vaccines, partly because of certain views about citizen liberty, and partly because they have culpably ignorant beliefs about the safety and efficacy of vaccines. Their collective aim is to distribute as much 'information' as they can to the general public about the dangers and inefficacy of vaccines, in a bid to erode public confidence in the use of vaccines. The campaign results in the production of huge numbers of false beliefs (and suspension of judgement) in the surrounding community about the efficacy of vaccines.

²²For detailed discussion of degrees of epistemic criticizability, see Boulton [forthcoming](#).

In addition to whatever epistemic failings an individual team member may commit in their participation in the campaign, individual team members also seem to bear some kind of responsibility for the larger epistemic harm caused by the campaign itself. After all, their individual efforts combine to collectively bring about those harms. But how exactly can the Causal Account make sense of this? Since there is such a large number of team members, the overall operation would have been successful even with just one fewer participant. So, in a sense, it is implausible that every individual team member really made a meaningful causal *difference* to the outcome of the campaign. In a nutshell, it seems relatively straightforward to come up with cases where a large number of people collectively bring about an epistemic harm – where each person seems in some sense complicit in that epistemic harm – but where the *absence* of any one of those people would not have made a difference to whether that harm came about.

At this point we have two main options. One approach is to simply adopt a highly pluralistic conception of (epistemic) complicity, accommodating causeless complicity in some other way than through a causal account. I have already said I want to remain neutral on this possibility. However, one potential worry about such an approach is that it leaves us with the task of explaining what these seemingly quite disparate phenomena have in common such that they all count as forms of *complicity*. Another approach is to accommodate ‘causeless complicity’ within a causal framework. Here I consider a few ways this might be done.

One way of addressing the issue of causeless complicity within a causal framework is by refining our understanding of ‘causal contributions’. On one approach, X counts as a causal contribution to Y just in case X is a necessary element of a sufficient set (NESS) of actual conditions of Y occurring (Hart and Honoré 1962). Consider a vote with four voters, A, B, C and D. All that is needed for a motion to pass is that three voters support it. Let’s say that on a given vote, all four voters support the motion. In this case, none of the individual votes is sufficient, nor necessary, for the motion to pass – after all, had any one of the individual voters decided not to support, the motion would still have had three supporters. We might take this to imply that none of the voters made a causal contribution to the passing of the motion. But that seems odd. To capture intuitions otherwise, we merely observe that each vote was a necessary element of a sufficient set of actual conditions on the motion passing. Had any one of them voted otherwise, the actual conditions leading to the motion passing would have been different.

Applied to the 10,000 anti-vaxxers case, the thought is as follows. Even though the disinformation campaign would have succeeded had any one of the distributors decided to stay home that day, it’s nevertheless true that each distributor is a necessary element of the sufficient set of actual conditions allowing the campaign to succeed. Had any of them stayed home that day, the specific set of sufficient conditions leading to the campaign’s success would have been different. In this sense, each individual distributor makes a causal contribution to the epistemic harm attending the success of the campaign.²³

²³I am not very convinced by this approach. It strikes me as assuming that the actual set of sufficient conditions for an outcome are what matters in a way that simply begs the question about whether a given individual’s role vis-à-vis a collective outcome should be understood as *causal* or not. If we assume that it should not, then it seems entirely fair to say of any given distributor that they did not really *causally*

Another approach is to distinguish between *definitely essential*, *potentially essential* and *inessential* contributions. A contribution is *definitely essential* to an outcome just in case there is no near-enough possible world in which the outcome could have occurred without the contribution. A contribution is *potentially essential* just in case there is at least one near-enough possible world in which the outcome could not have occurred without the contribution. A contribution is *inessential* just in case there is no near-enough possible world in which the outcome could *not* have occurred without the contribution. Definitely essential contributions are paradigm cases of ‘making a difference’, and as such are considered genuine causal contributions by all parties to the dispute. But under the right circumstances, we may also legitimately count potentially essential contributions as genuine causal contributions (Lepora and Goodin 2013: 61). Considered from the prospective perspective – that is, from the point of view of an agent deciding what to do, and where a given outcome is as yet unknown – the fact that it can be reasonably expected that one’s action might (in some near-enough possible world) end up being essential for an outcome seems crucial when it comes to questions of complicity. Of course, looked at retrospectively, such a contribution may turn out to have been inessential for the outcome. But that shouldn’t be enough to get an agent who saw (or should have seen) their own contribution as potentially essential off the hook. According to Lepora and Goodin, this is why potentially essential contributions are genuine causal contributions.

This approach might seem to confuse epistemic considerations with metaphysical ones. For example, it may sound odd to say that someone who merely potentially made an essential difference to outcome X actually causally contributed to X. Isn’t this a rather theoretically loaded sense of ‘causal contribution’? To put the worry another way, perhaps this approach begs the question against non-causal approaches to complicity. Perhaps we can all agree that the ‘potentially essential contributor’ to X bears some *responsibility* for X. But since causal accounts of complicity make causal contribution a necessary condition on (vicarious) responsibility, proponents of those accounts go further and maintain this entails that the agent made a causal contribution to X. But potential essentiality doesn’t *seem*, on independent grounds, like a causal contribution. So perhaps this is begging the question.

In response, a core move in causal accounts of complicity is to get away from participatory intentions and focus on facts about what an agent can *reasonably expect* (Gardner 2007; Lepora and Goodin 2013: 81–83; Petersson 2013). The notion of a causal contribution is crucial here. But it’s facts about what an agent can reasonably expect their causal contribution to be, as opposed to what it actually ends up being, that matter for an agent’s complicity. Does this ultimately amount to abandoning causal contribution as a minimal condition on complicity? I don’t think so. It simply amounts to taking a certain stance on the *perspective* from which we evaluate the causal relationship between the actions of a secondary agent and a principal harm or wrongdoing.

Two final points. First, by allowing certain kinds of *omissions* to count as causal contributions, we can capture a wider range of cases in which it may otherwise seem

contribute to the success of the campaign, precisely because, even though the actual set of sufficient conditions for the success of the campaign included them (as a necessary element), it could have *very easily* been some different set. That’s exactly why it can sound odd to say that what underwrites the agent’s role in the outcome is something about their causal contribution, as opposed to other factors, such as ones about their intentions and shared goals. See Petersson (2013) for further discussion and defence of this approach.

implausible to attribute a causal contribution to an agent. Omissions naturally count as causal contributions within the essential/potentially essential framework. For example, so long as there is at least one near-enough possible world in which the outcome could *not* have occurred without the omission, an omission counts as potentially essential to an outcome, and so a causal contribution. Second, many cases of agents seeming to make *no* causal contribution are really just cases of agents making a *small* contribution (to a larger whole that is itself significant). Returning to a version of the Dresden case, assume there was no specific threshold of bomb-number and density needed to cause a firestorm. Even if there would have been a lot of destruction with just 999 bombers dropping bombs, it's nevertheless true that there was a little bit more given that there were 1000. So, the 1000th bomber causally contributed to the precise amount of destruction that was caused in the actual bombing. Things get trickier when an agent's contribution to X lies beyond a certain minimal threshold on outcome X. But that is what notions like potential essentiality are designed to help with.

As I said at the outset: my aim is not so much defending the Causal Account against all objections, or possible alternative accounts of epistemic complicity. It is rather: (i) to explain what it takes for an account of complicity to be *epistemic*; and (ii) to demonstrate the theoretical usefulness of epistemic complicity. I've already offered my defence of how the Causal Account can explain cases of genuinely epistemic complicity. I turn now to aim (ii), focusing on the role of 'shared epistemic responsibility' in a socially extended conception of epistemic blameworthiness.

6. Shared vs. vicarious epistemic responsibility

A number of authors have recently argued that there is an important role for *shared* epistemic responsibility in our epistemic communities (cf. De Haan 2021; Fleisher and Šešelja 2023). For instance, Boyd Millar (2021) argues that individuals can be responsible not only for their own individual epistemic conduct, but can also be *jointly* responsible for the joint epistemic conduct of the collective of which they are a part. While Millar's joint epistemic responsibility is something over and above individual epistemic responsibility, it is still a kind of *direct* epistemic responsibility. In this section, I want to argue that the idea of vicarious epistemic responsibility – and more specifically, complicity – provides an additional lens for examining the phenomena Millar takes to motivate his account. In some cases, it may be a preferable lens.

Millar is interested in cases where individuals seem blameless for their own individual epistemic conduct, but nevertheless seem epistemically blameworthy for *something*. Consider Heinrich, 'who lives in an isolated European community a few hundred years ago where everyone believes there is such a thing as witchcraft' (Millar 2021: 502). Heinrich grew up hearing countless stories from authoritative sources about witches; he has never encountered undermining or rebutting evidence casting doubt on the stories. Perhaps it would be unfair to epistemically blame Heinrich for believing in witchcraft. After all, he grew up in an epistemic environment which renders it unreasonable to expect that he could have done otherwise.²⁴

²⁴For defences of the idea that Heinrich is not individually epistemically blameworthy, see Goldman (1988: 51–52), Russell (2001: 36), Nottelmann (2007: 209–11) and van Woudenberg (2009: 382). See also Levy (2018: Section 3) for a broadly similar view.

However, according to Millar, it still feels natural to expect something more from people in this sort of situation. While it may be true that no individual member of the community can be reasonably expected to investigate the veracity of the stories about witchcraft, the community as a whole has ample resources for conducting such an investigation. Moreover, doing so seems extremely important because the community's belief in witchcraft is used to justify violence against women in the community. 'So, given that the belief in witchcraft is universal in the community, and given the weakness of the testimonial evidence supporting this belief, the members of this community can reasonably be expected to take some joint action to seek additional evidence' (2021: 502). According to Millar, we can account for this in terms of shared epistemic responsibility (2021: 503). While no individual can reasonably be expected to believe otherwise than they do, it *is* reasonable to expect the members of this community to jointly do what they can to inquire together in such a way that the individual members of the community have a better chance of enlightening themselves. So, while the individuals are not individually epistemically blameworthy, they are *jointly* epistemically blameworthy, precisely to the extent that they do not.

This case has structural analogues in important present-day scenarios, from large-scale situations such as a modern democratic society's sub-optimal epistemic position with respect to climate change, to the operations of biases in hiring committees (Millar 2021: 504–5). According to Millar, situations abound in which individuals cannot reasonably be expected to do epistemically better on their own, but where they nevertheless can reasonably be expected to comply with a shared epistemic obligation to engage in more epistemically advantageous joint epistemic conduct.

I will stick with Miller's witchcraft case, since the point I want to make about it plausibly generalizes to a wide range of other cases. My primary concern is as follows. If it is unreasonable to expect of any individual in this community that they could have believed differently about witchcraft, it is hard to see how it is reasonable to expect that they would join forces with other members of the community to ameliorate the broader epistemic situation regarding witchcraft. After all, from their perspective, there is nothing to improve. The issue goes beyond simply pointing out that it's hard to see what specific epistemic fault is at play in someone like Heinrich's failure to comply with the putative shared epistemic obligation. Perhaps at least for some of the community members, it is reasonable to expect at least a *bit* of curiosity, or perhaps even suspension of judgement about witchcraft, such that, even if outright rejection of their current doxastic states about witchcraft cannot reasonably be expected, then at least a willingness to join forces with others and get clearer on the matter is. This appears to be Millar's take on the case, as when he writes: 'plausibly, even an ordinary person such as Heinrich ought to recognize that the available evidence – stories of extraordinary events that occurred in the distant past or in faraway places – is problematic' (502). The deeper issue, in my view, is that it's far from clear whether members of Heinrich's community can reasonably be expected to have *executed* such a joint enterprise. This is because the execution of such a joint inquiry requires more than a mere willingness to join forces on behalf of individual members of the community. Minimally, it would seem to require a considerable degree of *confidence* in the worth of the project, and mutual *agreement* about how to proceed. These things seem needed to ensure a minimal kind of cooperation that would partially constitute a 'joint' inquiry into the matter. But it seems unreasonable to expect such confidence and agreement in this sort of situation. Indeed, that just seems to be a constitutive feature of the unlucky epistemic environment Heinrich finds himself in. So, if the community's individuals are not epistemically blameworthy

for their own epistemic failings, it seems neither are they epistemically blameworthy in any shared or joint sense either.

All of this said, I agree with Miller that there is a residue of responsibility here. And this is where I think the Causal Account of Epistemic Complicity is useful. I have said that I do not think it reasonable to expect that any individual member of the community would have the kind of confidence in the worth of joint inquiry and mutual agreement about how to proceed, needed to ensure a minimal kind of cooperation constituting a joint inquiry into the veracity of beliefs in witchcraft. But I also agree with Millar that it may be reasonable to expect, of any given individual, some curiosity and perhaps even suspension of judgement about witchcraft. Notice that the latter may be enough to explain how individual members of the community can, to some extent, be guilty of *culpably ignorantly contributing* to the collective ignorance of the community regarding witchcraft. For instance, it is not at all difficult to imagine individuals ‘looking the other way’, so to speak, or ‘turning a blind eye’, when it comes to an honest confrontation with their doxastic attitudes about witchcraft. In this way, perhaps (at least some) members of Heinrich’s community culpably ignorantly contribute to the collective beliefs of the community, even if it seems a stretch too far to expect the confidence and agreement needed between individuals for a full-scale joint inquiry. So long as a given individual knows or is culpably ignorant of the fact that the community’s falsely believing in witchcraft would be an epistemically harmful thing, they may meet the conditions on epistemic complicity according to the Causal Account. So, if we are inclined to regard the *collective* as in some sense culpably epistemically failing in its epistemic practices and doxastic states about witchcraft, we can explain our inclination to not let individuals completely off the hook in terms of a certain degree of culpability for complicity in the collective’s principal epistemic failing.²⁵

The Causal Account of Epistemic Complicity so far leaves wide open questions about what a community like Heinrich’s ought to do. This may seem like a defect. However, the Causal Account has action-guiding resources in this context as well. Using the framework’s mechanisms for degrees of culpability for complicity, the more culpable a member of Heinrich’s community is for their complicity in the principal epistemic wrongdoing, the more we should expect them to contribute to a solution to the problem. For example, the community leaders and teachers are plausibly more culpable for their complicity in the principal epistemic wrongdoing than, say, a member of the community who plays no leadership role whatsoever. We might say their causal contributions to the collective’s ignorance are more *essential*, and more *central*, for example. The idea, then, is that this difference in culpability grounds a stronger obligation on those individuals to do something to contribute to the epistemic amelioration of the community. This leads me to an important caveat. It seems plausible that there may of course be certain shared epistemic obligations amongst such subsets of members of the community. So, we need not think of my proposal as entirely at odds with Millar’s. If anything, perhaps they complement one another.

Millar argues that any account of epistemic blameworthiness and responsibility that leaves out *shared* epistemic blameworthiness and responsibility is radically incomplete. He takes his cases to demonstrate this. I share his aim of expanding our understanding

²⁵Here I assume that if we take something like Kutz’s minimal normative approach to collectives (see fn. 18), it is quite plausible that the community constitutes a kind of principal epistemic wrongdoer, one capable of grounding claims about vicarious responsibility in individual members.

of the social dimension of our epistemic blaming practices. But given that complicity seems to do a better job of explaining at least some of the motivating cases than Millar's appeal to shared epistemic responsibility, I would propose that any discussion of epistemic blameworthiness that leaves out not just *shared* epistemic responsibility, but also *vicarious* epistemic responsibility, is radically incomplete.

7. Conclusion

Why does epistemic complicity matter? One reason complicity in the moral domain matters is because it helps us better understand what's problematic about excuses like, 'But I was just following orders!' or, 'But my actions make no real difference!'. One reason these excuses fall flat is because, despite however blameless these individuals are for what they directly do, they may nevertheless be implicated in – vicariously responsible for – the wrongs or harms of another agent or collective. I have argued that complicity can do something similar for us in the epistemic domain.

It does so primarily by allowing us to more fully explain what is objectionable about an important class of epistemic agent who, on an individual level, may be epistemically blameless or blameworthy to a relatively low degree, but whose relation to other individuals or collectives nevertheless makes them epistemically blameworthy to a higher degree. Take Carl the newscaster: in addition to being obviously morally problematic, there is something deeply epistemically problematic about his activities. However, this is also somewhat puzzling because, as we know, Carl's methods for constructing his monologues are in many ways permissible from an epistemic point of view. In an important sense, he's not stupid; he's just cynical: he proportions his *own beliefs* to the evidence, and manifests a great deal of epistemic rationality.²⁶ It is unclear how to explain the full scope of an attribution of *epistemic* responsibility and blameworthiness here. Unless we appeal to something like epistemic complicity. The Causal Account of Epistemic Complicity can help explain the full scope of such an individual's epistemic responsibility. Carl makes a significant causal contribution to the culpable epistemic failings of a large collective, and the epistemic dimension of that harm renders him an appropriate target of blame. Carl is epistemically complicit in the culpable epistemic failings of his audience.

Similar points can be made about other cases, such as those discussed in the shared epistemic responsibility literature – cases where whatever direct epistemic failings an individual has committed may seem relatively minimal (or not), but where those very failings function as mechanisms of vicarious epistemic responsibility, thus amplifying the agent's overall epistemic blameworthiness. Importantly, this idea may be capable of shedding light on the relationship between individuals and collectives in other wrongs and harms attending a wide range of epistemic and hermeneutical injustices. So long as the relevant phenomena involve individuals culpably ignorantly or knowingly contributing to larger-scale epistemic harms (while knowing or being culpably ignorant of the fact of those epistemic harms), we may be able to understand these phenomena in terms of epistemic complicity. We may then deploy the significant

²⁶Again, if we think about Carl's actual broadcasting of the monologues, and not just his construction of them, it seems we must also take into account the role of assertion in assessing his epistemic blameworthiness. But, again, my aim here is not to argue that someone like Carl is completely epistemically blameless for his direct epistemic conduct. Rather, it is to argue that focusing on his direct epistemic conduct is insufficient in important ways for capturing the full scope of his epistemic blameworthiness.

theoretical recourses of the Causal Account to further illuminate their role in our epistemic communities. I leave such an expanded discussion for another time.

For now, what seems clear is that my core examples in this paper are not niche. Arguably, they are among the most important ways that people can be epistemically blameworthy. The literature on epistemic blame has tended to focus on the blameworthiness of individuals for their own epistemic failings (Boult 2021; Brown 2020a, 2020b; Piovarchy 2021; Schmidt 2021). The focus is almost exclusively on things like individual epistemic vices, irrational beliefs and faulty inquiry. If anything like the foregoing is correct, this focus is highly disproportionate to the full range of ways people can be epistemically blameworthy in the real world.

References

- Bazargan-Forward S.** (2017). 'Complicity.' In M. Jankovic and K. Ludwig (eds), *Routledge Handbook on Collective Intentionality*, pp. 327–37. London: Routledge.
- Bennett C.** (2021). 'Complicity and Normative Control.' *The Monist* 104(2), 182–94.
- Bjornsson G.** (2020). 'Collective Responsibility and Collective Obligations Without Collective Moral Agents.' In S. Bazargan-Forward and D. Tollefsen (eds), *Handbook of Collective Responsibility*, pp. 127–41. New York: Routledge.
- Blomberg O.** (2023). 'How to be Morally Responsible for Another's Free Intentional Action.' *Journal of Ethics and Social Philosophy* 25(3), 545–79.
- Blomberg O. and Hindriks F.** (2020). 'Collective Responsibility and Acting Together.' In S. Bazargan-Forward and D. Tollefsen (eds), *The Routledge Handbook of Collective Responsibility*, pp. 142–54. New York: Routledge.
- Boult C.** (2021). 'There is a Distinctively Epistemic Kind of Blame.' *Philosophy and Phenomenological Research* 103(3), 518–34.
- Boult C.** (2023). 'The Significance of Epistemic Blame.' *Erkenntnis* 88(2), 807–28.
- Boult C.** (forthcoming). 'Degrees of Epistemic Criticizability.' *The Philosophical Quarterly*.
- Bratman M.** (1993). 'Shared Intentions.' *Ethics* 104, 97–113.
- Brown J.** (2020a). 'What is Epistemic Blame?' *Noûs* 54(2), 389–407.
- Brown J.** (2020b). 'Epistemically Blameworthy Belief.' *Philosophical Studies* 177(12), 3595–614.
- Cassam Q.** (2019). *Vices of the Mind: From the Intellectual to the Political*. Oxford: OUP.
- Collins S. and De Haan N.** (2021). 'Interconnected Blameworthiness.' *The Monist* 104(2), 195–209.
- De Haan N.** (2021). 'Collective Culpable Ignorance.' *Thought* 10(2), 99–108.
- Driver J.** (2015). 'Individual Consumption and Moral Complicity.' In B. Bramble and B. Fischer (eds), *The Moral Complexities of Eating Meat*, pp. 67–79. Oxford: Oxford University Press.
- Feinberg J.** (1966). 'Causing Voluntary Actions.' In W.H. Capitan and D.D. Merrill (eds), *Metaphysics and Explanation: Proceedings of the 1964 Oberlin Colloquium in Philosophy*, pp. 29–47. Pittsburgh: University of Pittsburgh Press.
- Fleisher W. and Šešelja D.** (2023). 'Responsibility for Collective Epistemic Harms.' *Philosophy of Science* 90, 1–20.
- French P.** (2016). 'Complicity: That Moral Monster, Troubling Matters.' *Criminal Law and Philosophy* 10, 575–89.
- Fricker M.** (2007). *Epistemic Injustice*. Oxford: OUP.
- Galvin R. and Harris J.** (2023). 'Causal Impotence and Complicity.' *Public Affairs Quarterly* 37, 47–63.
- Gardner J.** (2007). 'Complicity and Causality.' *Criminal Law and Philosophy* 1(2), 127–41.
- Gilbert M.** (2000). *Sociality and Responsibility: New Essays in Plural Subject Theory*. Oxford: Rowman & Littlefield.
- Gilbert M.** (2002). 'Collective Wrongdoing: Moral and Legal Responses.' *Social Theory and Practice* 28(1), 167–87.
- Glavaničová D. and Pascucci M.** (2022). 'Making Sense of Vicarious Responsibility: Moral Philosophy Meets Legal Theory.' *Erkenntnis* 89, 107–28.

- Goetze T. (2021). 'Moral Entanglement: Taking Responsibility and Vicarious Responsibility.' *The Monist* 104(2), 210–23.
- Goldman A. (1988). 'Strong and Weak Justification.' *Philosophical Perspectives* 2, 51–69.
- Hammerton M. (2018). 'Distinguishing Agent-Relativity from Agent-Neutrality.' *Australasian Journal of Philosophy* 97(2), 239–50.
- Hänel H.C. (2020). 'Hermeneutical Injustice, (Self)-Recognition, and Academia.' *Hypatia* 35, 1–19.
- Hart H.L.A. and Honoré T. (1962). *Causation in the Law*. Oxford: Clarendon Press.
- Hurka T. (2003). 'Moore in the Middle.' *Ethics* 113(3), 599–628.
- Jackson D.L. (2019). 'Date Rape: The Intractability of Hermeneutical Injustice.' In W. Teays (ed.), *Analyzing Violence Against Women*, pp. 39–50. New York: Springer.
- Kadish S. (1985). 'Complicity, Cause, and Blame: A Study in Interpretation of Doctrine.' *California Law Review* 73(2), 323.
- Katz C. (2023). 'Consumer Complicity and the Problem of Individual Causal Efficacy.' *Journal of Ethics and Social Philosophy* 24(2), 1–27.
- Kutz C. (2000). *Complicity: Ethics and Law for a Collective Age*. Cambridge: Cambridge University Press.
- Kutz C. (2007). 'Causeless Complicity.' *Criminal Law and Philosophy* 1(3), 289–305.
- Lawson B. (2013). 'Individual Complicity in Collective Wrongdoing.' *Ethical Theory and Moral Practice* 16(2), 227–43.
- Lepora C. and Goodin R. (2013). *On Complicity and Compromise*, Reprint ed. Oxford: Oxford University Press.
- Lepora C. and Goodin R. (2016). 'On Complicity and Compromise: Reply to Peter French and Steven Ratner.' *Criminal Law and Philosophy* 10, 591–602.
- Levy N. (2007). *Bad Beliefs*. Oxford: Oxford University Press.
- Levy N. (2018). 'Socializing Responsibility'. In K. Hutchison, C. Mackenzie and M. Oshana(eds), *Social Dimensions of Moral Responsibility*, pp. 185–205. New York, NY: Oxford University Press.
- Lewis H.D. (1948). 'Collective Responsibility.' *Philosophy* 84, 3–18.
- Medina J. (2012a). *The Epistemology of Resistance: Gender, Racial Oppression, Epistemic Injustice, and the Social Imagination*. Oxford: OUP.
- Medina J. (2012b). 'Hermeneutical Injustice and Polyphonic Contextualism: Social Silences and Shared Hermeneutical Responsibilities.' *Social Epistemology* 26, 201–20.
- Mellor R. (2021). 'Faces of Vicarious Responsibility.' *The Monist* 104(2), 238–50.
- Millar B. (2021). 'Shared Epistemic Responsibility.' *Episteme* 18(4), 493–506.
- Mitova V. (2022). 'The Collective Epistemic Reasons of Social-Identity Groups.' *The Asian Journal of Philosophy* 1(2), 1–20.
- Montmarquet J. (1992). 'Epistemic Virtue and Doxastic Responsibility.' *American Philosophical Quarterly* 29, 331–41.
- Moore M. (2007). 'Causing, Aiding, and the Superfluity of Accomplice Liability.' *University of Pennsylvania Law Review* 156(2), 395–452.
- Nagel T. (1986). *The View From Nowhere*. New York: Oxford University Press.
- Nottelmann N. (2007). *Blameworthy Belief: A Study in Epistemic Deontologism*. Dordrecht: Springer.
- Owens D. (2000). *Reason Without Freedom: The Problem of Epistemic Normativity*. New York: Routledge.
- Palermos O. (2022). 'Responsibility in Epistemic Collaborations.' *Philosophical Issues* 32, 335–50.
- Parfit D. (1984). *Reasons and Persons*. Oxford: Clarendon Press.
- Peels R. (2016). *Responsible Belief: A Theory in Ethics and Epistemology*. Oxford: OUP.
- Petersson B. (2013). 'Co-responsibility and Causal Involvement.' *Philosophia* 41, 847–66.
- Pettit P. (1987). 'Universality Without Utilitarianism.' *Mind* 72, 74–82.
- Piovarchy A. (2021). 'What Do We Want from a Theory of Epistemic Blame?' *Australasian Journal of Philosophy* 99(4), 791–805.
- Pohlhaus G. (2012). 'Relational Knowing and Epistemic Injustice: Toward a Theory of Willful Hermeneutical Ignorance.' *Hypatia* 27, 715–35.
- Portmore D. (2013). 'Agent-Relative vs. Agent Neutral.' In H. LaFollette (ed.), *International Encyclopedia of Ethics*, pp. 162–71. New Jersey: Wiley-Blackwell.
- Ridge M. (2005/2022). *Reasons for Action: Agent-Neutral vs. Agent-Relative*. Stanford Encyclopedia of Philosophy. At <http://plato.stanford.edu/archives/fall2008/entries/reasons-agent/>.

- Russell B.** (2001). 'Epistemic and Moral Duty.' In M. Steup (ed.), *Knowledge, Truth, and Duty: Essays on Epistemic Justification, Responsibility, and Virtue*, pp. 34–48. New York, NY: Oxford University Press.
- Scanlon T.M.** (1998). *What We Owe to Each Other*. Cambridge, MA: Harvard University Press.
- Scanlon T.M.** (2008). *Moral Dimensions: Permissibility, Meaning, Blame*. Cambridge, MA: Belknap Press.
- Schmidt S.** (2021). 'Epistemic Blame and the Normativity of Evidence.' *Erkenntnis* **89**(1), 1–24.
- Schmidt S. and Ernst G.** (2020). *The Ethics of Belief and Beyond: Understanding Mental Normativity*. Abingdon: Routledge.
- Schwenkenbecher A.** (2022). 'How We Fail to Know: Group-Based Ignorance and Collective Epistemic Obligations.' *Political Studies* **70**(4), 901–18.
- Shoemaker D.** (2011). 'Attributability, Answerability, and Accountability: Toward a Wider Theory of Moral Responsibility.' *Ethics* **121**, 602–32.
- Smith A.** (2008). 'Control, Responsibility, and Moral Assessment.' *Philosophical Studies* **138**, 367–92.
- Smith A.** (2012). 'Attributability, Answerability, and Accountability: In Defense of a Unified Account.' *Ethics* **122**, 575–89.
- Van Woudenberg R.** (2009). 'Responsible Belief and Our Social Institutions.' *Philosophy* **84**, 47–73.
- Zagzebski L.** (1996). *Virtues of the Mind*. Cambridge: CUP.
- Zakaras A.** (2018). *Complicity and Coercion*. Oxford: Oxford University Press.