Canadian Journal of Philosophy (2025), 1–15
doi:10.1017/can.2025.5

ARTICLE

Science Communication, Paternalism, and Spillovers

Hrishikesh Joshi

Department of Philosophy, University of Arizona, Tucson, AZ, USA
Email: hrishikeshjoshi@arizona.edu

Abstract

Epistemic paternalism involves interfering with the inquiry of others, without their consent, for their own epistemic good. Recently, such paternalism has been discussed as a method of getting the public to have more accurate views on important policy matters. Here, I discuss a novel problem for such paternalism—*epistemic spillovers*. The problem arises because what matters for rational belief is one's total evidence, and further, individual pieces of evidence can have complex interactions. Because of this, justified epistemic paternalism requires the would-be paternalist to be in an unusually strong epistemic position, one that most would-be paternalists are unlikely to meet.

Keywords: expertise; intellectual humility; paternalism; evidence; epistemic trespassing; spillover effects

1. Introduction

Old-fashioned paternalism involves interfering with someone's actions, or restricting their freedom of choice, for their own good. Thus, a state may mandate the wearing of seatbelts so as to protect people from serious injury or death in cases of car accidents. Recently, the notion of *epistemic* paternalism has attracted some interest. A common way to cash out this sort of paternalism is that it involves interfering with a person's inquiry without their consent for their own epistemic good (Ahlstrom-Vij, 2013; Jackson, 2022). Of course, what counts as interference with someone's inquiry can be difficult to determine in many cases. And people will disagree about what our epistemic good consists in, and how we might weigh different sorts of desiderata against each other—knowledge, true belief, understanding, and so on.

However, to fix ideas, an example from Goldman's (1991) original paper on the topic is illustrative. The central case there involves a judge blocking certain evidence from the jury so as to improve the quality of their deliberations and, ultimately, to get them to more reliably track the truth. It is useful here to compare the interpersonal norm that comes out in this case with intrapersonal norms regarding evidence. For one, it is widely accepted that an individual ought to base their beliefs on their *total* evidence, rather than a proper subset. Relatedly, an individual ought to gather all the evidence they can at negligible cost. However, in this case, the judge blocks certain evidence from the jury precisely to ensure that they deliberate based on a proper subset of the available evidence. The judge thus does not abide by the following putative norm:

> If an agent X is going to make a doxastic decision concerning question Q, and agent Y has control over the evidence that is provided to X, then, from a purely epistemic point of view, Y should make available to X all of the evidence relevant to Q which is (at negligible cost) within Y's control (Goldman, 1991, p. 114).

Some examples of the types of evidence that may be blocked in this way in the legal context include hearsay, details about past crimes, and testimonies regarding the defendant's character. In Goldman's discussion, the "epistemic good" of the jurors is cashed out in veritistic terms—the goal is to improve the likelihood of them getting at the truth of the matter. The jury is blocked from collecting evidence as it sees fit and, moreover, is not consulted on the matter. These features plausibly fit the interference and lack of consent conditions mentioned earlier.

Now, epistemic paternalism has attracted some recent interest in the context of scientific communication, especially with respect to politically polarized issues. The case of climate change is of special relevance in contemporary discussion. The problem is the following: while there is widespread scientific consensus that severe climate change is occurring and is attributable to human activity, a significant proportion of the public is skeptical of this claim (Anderson, 2011; Kahan, 2012). Because the reality of climate change is highly relevant to important human interests, some have argued that epistemic paternalism would be justified insofar as it would be effective in getting people to come to the right view on the issue (McKenna, 2020, 2023). In this general vein, it is thought that science communication might benefit from "marketing methods," designed to persuade people in ways congenial to their social and cultural values (Kahan & Braman, 2006; Kahan et al., 2011).

That said, not all uses of "marketing methods" need to count as epistemic paternalism, because for something to count as paternalism, it needs to satisfy the interference condition. Insofar as the marketing of this sort simply frames issues so that they connect with the values of the audience, it does not seem paternalistic. However, insofar as marketing interferes with the public's inquiry (say, if it suppresses evidence that is not congenial to the public's values), then it falls within the purview of the normative concerns related to epistemic paternalism.[1]

In this vein, John (2018) goes further to argue that we should not assume that putative scientific norms of openness, honesty, transparency, or sincerity are applicable in public communication, particularly on issues like climate change. Thus, scientists are justified in making somewhat misleading assertions, withholding certain findings, and so on, if doing so would increase the likelihood of more widespread public acceptance of the main conclusions of climate science. Of course, part of the rationale for these measures would be straightforwardly practical and non-paternalistic: having the public on board would aid effective policymaking with respect to climate change mitigation. However, we might also think there is a distinctively *epistemic* aim here. Getting some members of the public to have a more accurate picture of the world might have to involve interfering with their inquiry in certain ways, without their consent.

Now, old-fashioned paternalism faces familiar challenges. One such challenge that comes up in Mill's (1859) classic treatment of the issue in Chapter 4 of *On Liberty* is the following. People differ widely in their tastes, personalities, life experiences, dispositions, and so forth. Doing paternalism well will involve knowing the individual's specific dispositions and circumstances. However, the thought goes, would-be paternalists (e.g., legislators and authorities) will often lack access to these sorts of facts about the details of people's lives. And so, their interventions on our behalf are likely to be counterproductive.

In this article, I want to explore whether there might be an analogous reason that counts against epistemic paternalism. Now, at the outset, it might be thought that we simply lack the normative

[1]A tricky case here is thinking about a "marketing maneuver" that gets people to form a true belief *P*, without changing their evidential base (e.g., we can think of subliminal messaging that aims to change beliefs). It is not obvious how to think of this in terms of the subject's epistemic good. On the one hand, a true belief plausibly counts as an epistemic upgrade. But we also might want people's beliefs to be justified, that is, based in the right way, on the right sorts of evidence. Those sorts of epistemic goods would be missed by this type of marketing maneuver. What is "evidence" in this sense? For my purposes, I want to remain neutral with respect to different conceptions of evidence—however, in general terms, as Kelly (2014) puts it, "[e]vidence, whatever else it is, is the kind of thing which can make a difference to what one is justified in believing or (what is often, but not always, taken to be the same thing) what it is reasonable for one to believe."

https://doi.org/10.1017/can.2025.5 Published online by Cambridge University Press

authority to interfere with others' inquiry for their own epistemic good (Bullock, 2018). However, suppose that paternalist rejects this worry on the following grounds. There are some epistemic goods, it might be thought, that can outweigh sufficiently minor interferences with one's inquiry. Withholding a piece of evidence that is highly likely to mislead someone on a sufficiently important matter (from an epistemic perspective) might be worth the payoff. For a relatively minor infringement on epistemic autonomy, construed here as an agent's being able to conduct her inquiry as she sees fit, we secure an important epistemic good—say, knowledge or justified belief on some significant issue.

Thus, in what follows, I do not want to assume that there is some autonomy-based normative default against epistemic paternalism.[2] Rather, I want to argue here that even if we bracket autonomy-based considerations, epistemic paternalism will be difficult to justify for reasons analogous to the point made above against old-fashioned paternalism. My argument builds on the observation that the epistemic import of a new piece of evidence depends on one's total background evidence. This raises the possibility that paternalistic interventions geared at getting someone to form an accurate belief P might cause what I call *epistemic spillover* effects—unintended epistemic consequences relative to some other proposition Q. Avoiding such epistemic spillovers requires that the paternalist be in an unusually good epistemic position. She must know (or justifiably believe) not only P but must also be sufficiently cognizant of the total epistemic state of the individual on whose behalf she is conducting the paternalistic intervention. Only then can she be relatively sure that negative epistemic spillovers will not occur. And this, I argue, is often likely to be quite difficult.

Before proceeding, I want to fend off an initial worry having to do with the potential ubiquity of epistemic paternalism. It might be thought that in general, in almost any testimonial exchange, speakers must pick and choose what information to emphasize, what to simplify, and so on. And often, this is done with the epistemic interests of the hearer in mind—we do not want to overwhelm our listeners with the more insignificant or technical details for instance. Similarly, when scientists produce reports, they will want to highlight the most significant information in a way that is most digestible by their readers. They will not simply want to overload their audience with more technical details from the academic literature. However, of course, these practices are surely desirable and practically unavoidable.[3]

While these practices have the epistemic good of others in mind, they do not qualify as epistemic paternalism of the sort that is at issue here, in part because they do not meet the interference condition. As Ahlstrom-Vij (2013, p. 41) characterizes it, interfering with someone's inquiry involves constraining her "ability to access, collect, and evaluate information in whatever way she happens to see fit." This seems to be satisfied in Goldman's judge case, but not in the case where, for instance, someone summarizes the results of a study, highlighting the important points, while trying not to overburden readers.

A more subtle and interesting case is as follows. Suppose that the background data and literature in some domain is sufficiently technical or practically inaccessible to the broader public. Suppose further that the data/literature suggests two broad claims, P_A and P_B. However, to get the public to come to a particular conclusion, P_C (which, let us assume, is true), the researchers in this domain only present P_A. This seems to satisfy the interference condition in the way the action of Goldman's judge does.

What must be done to avoid epistemic paternalism here? A plausible first pass here seems to me to be that the researchers must offer a *representative summary* of the available evidence, such that the audience can then make up their own minds as to what that evidence supports on whatever range of questions they might be interested in.[4] This is not to say that P_B must be mentioned whenever P_A is in all contexts. Rather, more weakly, the thought is that P_B must be made sufficiently accessible to the public. That said, the notion of "interference with inquiry" is bound to be vague,

[2]Thanks to an anonymous referee for pressing me to make this clear.

[3]Thanks to the referees for raising this point.

[4]See Ballantyne (2015) for a detailed discussion of the notion of evidential representativeness in this sense.

https://doi.org/10.1017/can.2025.5 Published online by Cambridge University Press

and there will be gray areas where it is not clear whether interference is going on. For example, it is not clear what to say when P_B is stated in the report but is buried in footnote 38.

What is more, some authors have argued that even certain sorts of rational persuasion can count as paternalism. Tsai (2014) gives an example of a father who wants his daughter to go to law school rather than philosophy graduate school, and as a result, overwhelms her with reasons why she should study law. Here, he is motivated by distrust in his daughter's ability to make the right call. Tsai's thought here is that the father can intrude upon his daughter's deliberation, even if he only uses rational persuasion—that is, gives reasons and arguments as opposed to employing other nonrational techniques.[5]

In what follows, however, I will focus on the central sort of case that Goldman begins with—that is, of someone withholding evidence for the sake of another's epistemic good. This is for two reasons. First, it will help to best illustrate the idea of epistemic spillovers that I am interested in here. Second, withholding evidence is connected to the issue of censorship. Some forms of censorship can constitute a sort of epistemic paternalism—in particular, where the censor's aim is to block certain evidence from reaching an audience for the audience's own epistemic good. However, censors can be motivated by concerns apart from an audience's epistemic good—for instance, their own material interests or harm to third parties (Messina, 2023).

2. Epistemic Spillovers

In economics, spillovers (also called "externalities") are the effects (usually negative) of some economic transaction on third parties. From the perspective of social welfare, these effects are important to take into account. If negative spillover effects are strong enough, we might say there is a significant "market failure"—that is, the inefficient allocation of resources via market processes, given a particular set of property rights.

To illustrate, trade is usually a benefit to both parties—so long as they are informed, not coerced, and so forth. Thus, consider the buying of a laptop computer. Here, money is exchanged for the computer, and by each party's lights, there is a net benefit. The buyer sees the new computer as more valuable to her than, say, the 1000 dollars. And the seller sees the 1000 dollars as more valuable to them than the computer. So, both parties can be said to win here. However, this is only part of the picture from the perspective of social welfare. For, suppose that manufacturing the computer involves creating harmful chemical byproducts. And suppose the manufacturer dumps these into a nearby river. Now, there are negative spillover effects. Among other things, this dumping may be harmful to fishermen who will face reduced stocks, to people who might rely on the river's water for household purposes, and of course to the environment in general, which many of us want to preserve. Effective policymaking thus involves mitigating such externalities, for example, by placing an appropriate "Pigouvian" tax on pollution, banning the use of especially harmful chemicals, and so on.[6]

In a similar vein, I argue, there can be *epistemic* spillover effects.[7] Here is a stylized example to bring this out. Mustard believes that Wadsworth killed Green. Let us suppose that Mustard's belief

[5]Another sort of case involves *giving* evidence to people without their consent (Bullock, 2016, 2018). Thus, we can imagine a doctor disclosing information about a patient's medical condition despite her stating that she does not want to know about it. Paternalistic interventions of this sort thus need not involve withholding evidence—in fact, just the opposite. However, for the purposes of this article, I focus on cases of withholding—in part because they are the more common case, and in part because they bring out the problem of spillovers.

[6]For the classic discussion of this point, see Pigou (1920).

[7]Recently, the term "epistemic spillover" has been used to describe a distinct phenomenon. In an influential article, Marks et al. (2019) find that political partisans tend to inflate (deflate) the credibility of testimony from ingroup (outgroup) members even on subject matters having nothing to do with politics. In particular, they report that learning that someone is politically like-minded led participants in the study to take that person's judgment more seriously even on something totally unrelated to politics—in this case, a geometric shape recognition task. So, in this sense, our tendency to defer to in-group members on politics "spills over" to nonpolitical domains. However, this is distinct from the sense in which I am using the term "epistemic spillover." In my usage here, it refers to the epistemic analog of the concept of spillover (or externality) in economics.

https://doi.org/10.1017/can.2025.5 Published online by Cambridge University Press

is justified—he is appropriately responding to the total evidence he has. Furthermore, let us even say Mustard *knows* that Wadsworth committed the murder—Wadsworth did, in fact, do it. Now, suppose that here are the relevant pieces of evidence that Mustard possesses, as far as this case is concerned: (1) Wadsworth's fingerprints were on the murder weapon; (2) Wadsworth had a strong motive to kill Green—Green had recently discovered Wadsworth's insurance fraud scheme and was about to go to the police; (3) Wadsworth has no alibi; (4) Scarlett was seen leaving the mansion the evening that Green was killed; and finally, (5) Scarlett and Green had been married and had gone through a contentious divorce.

This above evidence is *mixed*, as evidence often is. Some of it points to Wadsworth, some of it to Scarlett. However, let us stipulate that an appropriate weighting of the evidence suggests that Wadsworth did it, not Scarlett. Now, Mustard is trying to convince Peacock that Wadsworth did it. Mustard knows that Peacock knows (5), and he thinks this prejudices Peacock unfairly against Scarlett. From Mustard's perspective, Peacock is prone to give (5) more weight than it deserves. Because of this, Mustard, in trying to convince Peacock that Wadsworth did it, intentionally withholds (4), which Peacock does not know. On the other hand, he tells her about (1), (2), and (3), and thereby succeeds in convincing Peacock that Wadsworth did, in fact, do it. In withholding evidence for the sake of Peacock's own epistemic good, Mustard acts like Goldman's judge from Section 1.

So far so good—Mustard has gotten Peacock to believe a true proposition, namely that Wadsworth killed Green. From a certain perspective, this is an epistemic upgrade for Peacock. Though, we may quibble about the details. Perhaps due to the omission of (4), she lacks the full picture, and maybe she does not achieve a full understanding with respect to how and why the crime occurred. However, maybe this is okay. I lack anywhere close to full understanding of why and how climate change is occurring, but I take myself to know *that* it is occurring. And this latter knowledge is an epistemic upgrade (and indeed, a significant one) relative to my being agnostic on the issue or believing that climate change is not occurring.

However, let us add now to the above case. There are further pieces of evidence that Peacock is privy to, which are not on Mustard's radar: (6) Scarlett works for Big Insurance and (7) Scarlett has significant gambling debt that she needs to pay off soon. Now, (6) and (7), let us suppose, are not by themselves, enough to point to conspiracy on Scarlett's behalf with Wadsworth's insurance fraud. However, when (4) is added to the picture, especially given the timing of Green's death, it strongly suggests that Scarlett is complicit in the insurance fraud. Had Mustard not withheld (4) from Peacock, she would have formed an important true belief: namely that Scarlett is in on Wadsworth's insurance fraud scheme.

Thus, Mustard's paternalistic intervention—that is, the withholding of a piece of evidence that he thinks will mislead Peacock regarding the murder—has a negative epistemic spillover. It has the effect of preventing Peacock from forming a significant true belief constituting knowledge. When viewed as a whole, it is not obvious that Peacock gets an epistemic upgrade due to Mustard's paternalism. Depending on whether and to what extent Peacock might have overweighted (5), she may well, on the whole, have suffered an epistemic *downgrade.*

Suppose that Peacock would have overweighted (5), taken in conjunction with (4), to such an extent that she would have believed Scarlett killed Green, not Wadsworth. In this case, she comes to a false belief regarding the murder, but also comes to know that Scarlett was in on the insurance scam. How to balance these against each other for an overall assessment depends on our epistemic axiology—how we weigh the good of true beliefs against the bad of false ones. It also depends on which beliefs we consider as more "important," from an epistemic perspective. (Presumably, my having a true belief about climate change is more important, even in a purely epistemic sense, than my having a true belief about how many blades of grass there are in my backyard.) But suppose that Peacock would have overweighted (5), but not to such an extent that she would have been led to believe Scarlett killed Green—she would have perhaps suspended judgment or believed it was likely Wadsworth but with less confidence than appropriate. Here, it seems that Mustard's paternalistic intervention plausibly leads to an epistemic downgrade for Peacock.

https://doi.org/10.1017/can.2025.5 Published online by Cambridge University Press

Note here that this is an *epistemic* reason to avoid epistemic paternalism. It is epistemic because the intervention has yielded an outcome that has downsides with regard to Peacock's epistemic good. That said, it does not follow that the intervention is not epistemically justified, all things considered.[8] For it may be thought that Peacock's coming to the true belief about the murder is more important than her coming to the true belief about the insurance fraud. The point is, however, that the significant possibility of an epistemic spillover provides *pro tanto* epistemic reason to avoid paternalism. In Section 3, I will present a *moral* reason to avoid epistemic paternalism.[9]

Before proceeding, I want to clarify one important difference between spillovers in the traditional economic sense and the sense I mean here. Spillovers in the economic sense are costs (or benefits) that accrue to *third parties*. You, as a buyer, and the seller of a computer both benefit from the transaction, but a third party—say, the fisherman on the nearby river—faces a cost due to the factory's pollution. In the sort of case, I am interested in, the (epistemic) cost does not accrue, in the first instance, to another *party* but rather with respect to an agent's doxastic attitudes with respect to another *proposition* (or propositions). Although of course, this is not to say such epistemic spillovers might not affect other individuals. Perhaps Peacock would have told others about Scarlett's complicity with Wadsworth's insurance fraud, and this would have led to an epistemic upgrade for them.

In this sense, epistemically paternalistic interferences can affect the rest of the community as well—indeed, as Clifford (1877), p. 292) emphasized, perhaps "no one man's belief is in any case a private matter which concerns himself alone." From a social epistemic perspective, our epistemic fate is bound up with that of others—thus, an epistemic spillover that affects one person is unlikely to be sequestered there, as it were. Nonetheless, on the way I am defining an epistemic spillover here, the key point is it *need not* affect a third party but might instead have a negative effect on some other doxastic state(s) of the individual on whose behalf the paternalistic intervention is conducted.

3. What Does it Take to Do Paternalism Well?

The above case brings out the fact that a new piece of evidence can have complex interactions with one's background total evidence. Thus, blocking a piece of evidence may have unforeseen epistemic consequences regarding some other proposition *Q*, even if it leads to an epistemic upgrade with respect to *P*—which is the focus of the paternalist's intervention. Indeed, there can also be cases where the *same* evidence supports *P* relative to my total background evidence whereas it supports not-*P* relative to yours (Kelly, 2008a). Insofar as disagreements are driven by differences in background evidence in this way, successful paternalistic intervention is likely to be difficult (Davies, 2022).

So, what does it take to do epistemic paternalism well—so that it avoids or minimizes negative epistemic spillovers? First, of course, the paternalist must have the right intentions and an adequate conception of what constitutes the epistemic good of others. She must also be reasonably good at evaluating the total evidence she has, as it bears on the target proposition(s) regarding which she intends to interfere with others' inquiry. Thus, if I am bad at evaluating evidence regarding economic data and policy, say, I should not engage in epistemic paternalism with respect to this area. Likewise, if Mustard is not good at analyzing evidence pertaining to criminal activity, he should not act paternalistically toward Peacock. What this suggests is that for someone to be justified in being epistemically paternalistic, they must have an adequate level of expertise about the subject matter pertaining to the target proposition(s).

[8]"Epistemic justification" in this sense is different from the sense typically used in the broader literature. As Jackson (2022, p. 144) notes, the term used in this context "doesn't pick out the thing that turns true unGettiered belief into knowledge. Here, 'justification' indicates when a practice, on balance, promotes epistemic goods."

[9]Thanks to an anonymous referee for asking me to clarify this distinction.

https://doi.org/10.1017/can.2025.5 Published online by Cambridge University Press

However, crucially, this is not enough. The paternalist must also have an adequate sense of the total evidence possessed by the target individual(s). This is what goes wrong in the case of the preceding section. Even if Mustard is an expert crime-solver, he does not have adequate access to Peacock's total evidence. This makes it hard to carry out the paternalistic intervention without unintended and unforeseen spillover effects. That is because pieces of evidence can have complex interactions with one another.

The problem would be less severe (though it still would not disappear) to the extent that evidence worked in a uniform, predictable way. Thus, suppose that as a general matter, each piece of evidence had equal weight, and either counted for, against, or was neutral with respect to any target proposition. Weighing evidence would then look like adding unit weights on either side of a scale and seeing where the balance tips. Here, it might sometimes be the case that the paternalist has so much evidence on one side of the scale that he may reasonably block inquiry, especially if he has good reason to think that the target individual is prone to errors of reasoning. Even here though, the paternalist must have a fairly good sense of the epistemic landscape surrounding the target proposition and be able to anticipate that significant negative spillovers would not occur.

There are, then, two broad sets of challenges that face the epistemic paternalist. First, he must have adequate expertise regarding the domain of the target proposition(s) but also regarding the domains where spillover effects might occur. Second, he must have an adequate sense of the target individual's total evidence so that he can be reasonably sure that negative epistemic spillovers will likely not occur. In the next two sections, I offer some reasons to think these conditions will often be difficult to meet.

Before proceeding, though, I want to highlight an important potential objection and explore how it might be addressed. In general, in the economic case, spillovers can be *positive* as well as negative. A beautiful building or mural, for example, has positive effects on others besides the individuals who may be parties to the relevant transaction. So, when it comes to epistemic spillovers, why not think that there can be positive spillovers of this kind? And further, it may be thought, if we are not assuming an autonomy-based default presumption against paternalism, then the justification for any particular instance of epistemic paternalism might depend on a sort of epistemic cost–benefit analysis. And there may be nothing useful to say in general about which way these analyses will tend to point.

In response, I want to note two things. First, to the extent that we think people are generally reliable in recognizing the import of new evidence, we should expect the spillovers to be negative. Note that this need not mean people are *always* reliable in this way with respect to every question. Indeed, this is the necessary assumption for getting the desirability of epistemic paternalism off the ground. However, the thought here is that people will be, for the most part, good enough, in general, at telling where the evidence points. Thus, in the case above, even if Peacock is prone to over-weighting (5), she is good enough at processing new evidence that she can see (1)–(7) as pointing to Scarlett's being in on the insurance fraud.

On the face of it, denying the above general claim seems to invite an unrealistic and unduly pessimistic view of our epistemic capabilities. In general, most people can see, for example, that the fact that the streets are wet is evidence that it rained, the fact that *X* is leading *Y* by 20 points in the polls is evidence that *X* will win the election, the fact that someone is crying is evidence they are upset, and so on. If we were *so terrible* at seeing where the evidence points, we would not have made it this far as a species. In this vein, Mercier (2020) has argued that we are reasonably vigilant when it comes to acquiring and incorporating new information in deciding what to believe. Of course, we have several well-documented biases and make use of rough and ready heuristics (Kahneman, 2011), but even these can be thought of as "optimally irrational" (Page, 2022)—that is, as trying to make the best use of limited cognitive resources. An analogy with our visual system might be apt here. We know that we are prone to various optical illusions. However, it does not follow that our visual capabilities are not, for the most part, highly reliable in our typical environment. The upshot

https://doi.org/10.1017/can.2025.5 Published online by Cambridge University Press

then is that insofar as we are generally good at seeing where new evidence points, the spillover effects of epistemic paternalism are likely to be negative.

I want to note one limit of this argument, however. If the paternalist has high enough expertise with respect to some domain *D* and if it can be reasonably predicted that the spillover effects will only affect claims within *D*, then epistemic paternalism might be carried out without significant risk of negative spillovers. However, when it comes to many policy-relevant matters, the spillovers might not be able to be "contained" within a specific domain of expertise in this way. I will say more about this in Section 5.

There is a second, *moral* reason to be especially concerned with negative epistemic spillovers, and this has to do with responsibility for the resulting harms. Suppose I withhold some evidence from you and, as a result, you (rationally) choose a course of action φ rather than ψ. But were you to have the evidence I withheld, the right course of action would have been to ψ. It seems that I am responsible, at least partially, for the harms that result from your φ-ing rather than ψ-ing. To go back to the earlier example, because Mustard holds (4) from Peacock, she does not have enough to go on to think that Scarlett was in on the insurance fraud. Suppose that as a result, she does not bring this up with the authorities, who are led to convict the innocent Mrs. White, given the evidence they have. Mustard seems at least partially responsible for this resulting injustice.

This asymmetry seems to mirror how we intuitively treat cases of negative spillovers. If someone paints a beautiful mural but as a result pollutes a nearby lake with harmful paint chemicals, they seem to bear the responsibility for the negative spillover in a way that is not canceled out by the positive effects of the mural. That is, they do not get off the hook simply by pointing to the mural. Indeed, this asymmetry seems to be reflected in our various legal and regulatory systems surrounding environmental regulation. Moreover, in this case, it is plausible that if a course of action carries a significant risk of such negative spillovers, that provides a strong reason against avoiding that action. For example, if the paint may or may not be harmful to the lake's ecosystem, the precautionary principle suggests that the paint should not be dumped there. Analogously, it is plausible that if there is a significant risk of epistemic spillovers, especially with regard to practically relevant propositions, that provides a strong reason against epistemic paternalism.

4. The "Fog of Debate"

Military commanders refer to the "fog of war" to describe uncertainty regarding the opposing side's position. Especially prior to modern communications and imaging techniques, it was very difficult to know about the enemy's movements, their weaponry, the terrain, and so on. Recently, Ballantyne (2021) has argued that efforts to persuade someone else of our position on an issue face similar challenges—we often operate within the context of a "fog of debate." For Ballantyne, there are four "clarity conditions," which may or may not hold within an argumentative context—to the extent they do not hold, we operate within a fog of debate. These conditions (Ballantyne, 2021, 94) are as follows:

> *Standpoint*: We accurately estimate how our audience thinks about an issue before we share our argument.
>
> *Comprehension*: Our audience understands our argument.
>
> *Force*: Our audience is justified to accept the argument's conclusion after understanding our argument.
>
> *Feedback*: We accurately interpret our audience's behavioral reactions to our argument.

Ballantyne's focus here is the context of argumentative persuasion, where we attempt to persuade someone of a position we hold, and they do not. The fog of debate can make our attempts to

https://doi.org/10.1017/can.2025.5 Published online by Cambridge University Press

persuade fall flat or even backfire. For example, to the extent that *Standpoint* fails to hold, we might attribute to our opponents a more extreme (or silly) position than what they, in fact, believe. As a result, we might end up "straw manning" the opposition, which may lead them to (perhaps rightly) take us less seriously going forward. *Feedback* can also fail to hold for many reasons and frustrate our attempts to persuade. For instance, sometimes norms of politeness and decorum might lead people to display faux agreement and we may end up thinking we have given a persuasive argument when, in fact, we have not.

Epistemic paternalism faces similar challenges. As I have been framing it, successful intervention requires the paternalist to have an adequate grasp of the total evidence that the target individual possesses. *Standpoint* and *Force* are related to this condition. As far as *Standpoint* is concerned, accurately estimating our audience's position on an issue means understanding not only their doxastic attitude toward the target proposition, but also the *reasons why* they believe what they do. *Force* also depends on the total background evidence the audience has. Indeed, on some accounts of rational polarization, two individuals may move *apart* on their attitudes toward some proposition *P* (so that one becomes more confident in *P* and the other becomes more confident in not-*P*) even after encountering the same mixed evidence.[10] The important point for our purposes is that paternalistic interventions also operate from behind a fog—even assuming no irrationality on the part of the paternalist, she must be able to discern the target individual's relevant total evidence.

Empirical evidence on this issue, particularly when it comes to political disagreement, is not encouraging news for paternalists. It turns out that partisans are often mistaken about what the other side believes—and particularly, they tend to attribute to the other side more extreme positions than they actually hold. For example, in the US context, partisans tend to value democratic principles (such as free and fair elections) but significantly underestimate the extent to which out-party members value those principles (Pasek et al., 2022)—potentially driving democratic backsliding (Braley et al., 2023). Further, a range of studies brings out the prevalence of "false polarization"—that is, a misperception of the distance between in-party and out-party members on various policy issues such as taxation, affirmative action, and so on (Blatz & Mercier, 2017; Levendusky & Malhotra, 2016; Sherman et al., 2003). Interestingly, individuals who are more politically active and more involved in trying to persuade others are more likely to misperceive what their political opponents think (Westfall, Van Boven, Chambers, & Judd, 2015). Thus, perhaps ironically, would-be epistemic paternalists, at least when it comes to politically charged topics, are less likely to have an accurate sense of the target individual's total evidence, in the way that is required to do paternalism well.

5. Expertise and the Problem of Containment

Successful epistemic paternalism requires *expertise.* Goldman is keen to emphasize this in his original article. He says, "To justify any particular instance of such paternalism, involving a particular controller, we must have grounds for taking that agent to be an expert" (Goldman, 1991, p. 128). For Goldman, the problem then becomes primarily one of *identifying* the relevant experts. Although there can be some difficult cases where experts disagree (Goldman, 2001), in many cases finding the relevant experts will be straightforward enough, he argues.

[10]See Kelly (2008a) for an influential discussion of this phenomenon. For Kelly, people can polarize in this way without any irrationality involved. Where real-world cases are involved, there are alternative explanations, notably based on the cultural cognition hypothesis (Kahan et al., 2017). The basic thought here is that we hold certain beliefs for reasons of identity—because we belong to certain groups, and it is an important marker of those groups to have certain beliefs. Thus, we are prone to motivated reasoning when it comes to evidence that challenges those beliefs. For a recent discussion of how the data might be captured with an alternative hypothesis—on which differences in background beliefs explain the polarization, rather than cultural cognition—see Davies (2022).

https://doi.org/10.1017/can.2025.5 Published online by Cambridge University Press

On a very simple model, it would seem the task of doing epistemic paternalism well proceeds as follows. There is some proposition *P* that experts agree is true but that a significant proportion of the public believes is false, or they are agnostic about it. If these latter beliefs are recalcitrant enough, epistemic paternalism is justified. This will involve restraining the inquiry of those who hold such beliefs. This interference may involve a range of things, but one main method is blocking evidence via various means of "communication control," as Goldman puts it. But who is to be charged with this control? The answer is simple enough: whoever constitutes the group of experts with respect to the subject matter that *P* belongs to. If *P* is a historical claim, then the relevant group is historians; if *P* is a climatological claim, then the group is climate scientists, and so on.[11]

The problem though is that fields of genuine expertise are *narrow*. This is particularly true in a time like ours, where the sciences have developed a high degree of division of labor so that expertise requires (hyper)specialization (cf. Hardwig, 1991). Thus, genuine experts are experts with respect to some highly specialized domains.

On the other hand, many issues of practical policy relevance do not fall within the somewhat artificial bounds of disciplinary demarcation. How best to tackle climate change? This question is essentially multidisciplinary—the considerations relevant to this question fall within the purview of economics, other social sciences, chemistry, physics, geology, and so on. Similar points hold regarding a host of other policy questions—pandemic responses, geopolitical relations, immigration policy, and so on.

This puts the would-be paternalist in a dilemma. On the one hand, a disciplinary expert can, we may reasonably assume, get the target individual to come to an accurate view about some proposition *P*, which falls within that expert's purview. But precisely because paternalism involves *interfering* with the individual's inquiry—by withholding evidence, misrepresenting the level of scientific certainty, strategic oversimplification, censorship, and so on—it poses a significant risk of epistemic spillovers with respect to complex policy questions.

The structure of this problem is analogous to the case of Mustard and Peacock from Section 2. Suppose there is a piece of evidence *E* that is relevant to whether *P*. Suppose that *E* is weak evidence *against P*, which is a true proposition that the expert knows. However, the expert worries that the target individual will overweight *E* and this will cause her to believe not-*P*, or perhaps suspend judgment. (Or more weakly, it might cause her to have more doubts about *P* than is warranted.) As far as the target individual's epistemic position with respect to *P* is concerned, the paternalistic intervention (i.e., the withholding of *E*) yields an upgrade.

However, given that complex policy questions draw on evidence from separate fields of inquiry, *E* is relevant to the question whether *Q*. Suppose *Q* is part of a subject matter that is not within the paternalistic expert's purview. Because *E* is relevant to whether *Q*, the target individual now forms an inaccurate view about *Q*, depending on just how much evidential import *E* has in this domain. We can call this the problem of containment—paternalistic interventions with respect to one domain risk epistemic spillovers into other domains.

Let me use a (highly simplified) example to illustrate. (The reader may feel free to use another case to fix ideas, if they do not think this example is plausible.) Suppose a person is skeptical about the severity and cause of climate change. Assume this person is on the conservative side of the political spectrum, and in the background is worried that climate activism is a ruse to set more regulations, get more state control over the economy, and so forth. An effective persuasive strategy with respect to this person might be to highlight the possibility of new "green jobs," innovation,

[11]It is interesting in this vein to think about the requirement of expertise as it applies in Goldman's judge case. On the one hand, it might be thought that the judge lacks expertise on general questions about human psychology and how individuals update on new evidence. However, it might be argued that the judge is an expert when it comes to the dynamics of trials, as well as the sorts of information that enable juries to more reliably track the truth about the case at hand. Alternatively, it might be claimed that the locus of the relevant expertise is the group responsible for developing and implementing these rules of evidence. For my purposes here, I wish to remain neutral on this more substantive issue.

https://doi.org/10.1017/can.2025.5 Published online by Cambridge University Press

increased exports, and so forth, that climate legislation might bring about (cf. Kahan et al., 2011).[12] This might make her more inclined to accept the proposition that climate change is severe and caused by human activity.

However, suppose the reality is that climate regulation is bound to include trade-offs. That is, we might have to forgo some by way of economic growth and current standard of living to ensure a safer planet for future generations. Depending on the persuasive strategy we take, this person might come to have a *false* view of this question—that is, she might come to think there is no tradeoff, or that the tradeoff is small. So, we get the person coming to believe one true proposition and one false proposition. Her coming to believe this false proposition is an epistemic spillover. Now it is not obvious whether the total change in her attitudes is indeed an epistemic upgrade.[13]

Efforts to deal with the problem of containment bring us to the other horn of the dilemma. For the expert to withhold (or suppress) *E* while being reasonably confident that it will not result in negative spillovers, she needs to make judgments about how *E* affects inquiry into other domains. But in doing so, the expert must depart from the boundaries of her subject matter expertise—in other words, she must engage in what Ballantyne (2019) has called *epistemic trespassing*.

Trespassing of this kind is problematic for both epistemic and ethical reasons. The epistemic problem is that the expert is making calls with respect to domains where she lacks the relevant expertise. This displays the epistemic vice of hubris. Even leading experts within a particular domain may simply lack the tools and background knowledge required to properly evaluate claims within other, even nearby, domains—as the case of Linus Pauling vividly brings out. Pauling was an expert if there ever was one, having won the Nobel Prize in chemistry and being celebrated as one of the top scientists of the twentieth century. However, he also claimed that high doses of Vitamin C can be a cure for various ailments, including cancer, which turned out to be unsupported by the evidence. In making these strong claims about Vitamin C, we might say he trespassed, by stepping outside the bounds of his expertise—from chemistry to medical science.[14]

As DiPaolo (2022) has pointed out, there are ethical problems with epistemic trespassing as well. When an expert addresses a novice, the latter typically has to surrender their epistemic autonomy. Thus, when my doctor prescribes a medicine to me, the appropriate response from my end is usually to take it for granted that the drug is appropriate, rather than trying to figure things out myself. However, this sort of surrendering of autonomy puts novices in a *vulnerable* position—outsourcing our beliefs in this way opens us up to a range of epistemic and practical harms. When experts pose as such in areas in which they lack disciplinary expertise, they are neglectfully abusing their authority.

The way out of this dilemma is to appeal to experts who can span a range of fields. Thus, perhaps there are genuine experts on public policy or geopolitics, or economics *simpliciter* rather than a particular subfield of research in these areas, like game theory or market microstructure. If there are

[12]The cited paper itself does not quite defend measures that rise to the level of epistemic paternalism. However, we could imagine policy proposals that build on the cultural cognition hypothesis to withhold information (or interfere with inquiry in some other way) as a means of getting the broader public to accept the scientific consensus on climate change.

[13]The recent COVID-19 pandemic might also furnish some examples of epistemic spillovers, and costly ones at that. Consider the claim that *COVID-19 is a severe disease and calls for significant policy intervention.* I take this claim to be true and justified, indexed to 2020. It seems right to conclude that experts engaged in various forms of epistemic paternalism to get people to have this true belief, in part by not communicating transparently (cf. Nguyen, 2024). However, on a plausible interpretation (though not the *only* plausible interpretation) of things, an epistemic spillover was created—namely, that *elementary school closures are justified* (for roughly the time that they were). For some empirical support of this general claim, see Graso et al. (2022), who found that greater belief in COVID-19 science and scientists also predicted pandemic mitigation authoritarianism. Dasgupta (2022) has recently argued that elementary school closure policies were on the net very harmful and unjustified relative to the total available evidence at the time. Furthermore, Dasgupta argues that closures were particularly harmful to disadvantaged students and families. If all this is right, the case illustrates the practical (not just epistemic) costs of spillovers. But again, the argument of this article does not rest on this particular case, which is only meant for illustration—as need be, the reader may substitute other cases they find more compelling.

[14]For a helpful discussion of the Pauling case and its epistemic implications, see Mercier and Sperber (2017).

https://doi.org/10.1017/can.2025.5 Published online by Cambridge University Press

such experts, then perhaps they can legitimately engage in epistemic paternalism without the risk of spillovers *or* trespassing.

However, the prospects for such expertise are dim. Putative general experts of this kind—or "pundits"—have a poor track record at prediction. Notably, Tetlock's (2005) extensive studies on the topic revealed that experts in geopolitics, for example, had no better success at prediction than lay people, or even dart-throwing chimpanzees. Or consider the financial crisis of 2008. People like Alan Greenspan or Ben Bernanke—experts on economics if there are any—failed to see it coming (Applebaum, 2019). The common wisdom had been that there was no "housing bubble" in the US market. Tetlock's work further finds that the people who were better at predicting such types of events typically were more provisional about their predictions and were loath to give yes or no answers (Tetlock & Gardner, 2015). In other words, better predictors are relatively more intellectually humble and open-minded.

What all this seems to suggest is the following. When it comes to policy, the evidential landscape is very complex and it is very hard for putative experts to predict what will happen. The best we can do is to make provisional claims and be open to new evidence as it comes along. Engaging in epistemic paternalism cuts against these points for two chief reasons. First, the gathering of evidence requires inquiry and interference with this inquiry is likely to frustrate the aim of gathering good evidence—especially in cases where such evidence might be discovered by novices or dissenting experts.[15] Second, epistemic paternalism is less likely to be justified the less certain one is about the target proposition. If I have no clue whether it will rain in 7 days, presumably I ought to be less disposed to interfere with others' inquiry about this proposition. However, as mentioned above, the individuals who were best able to make relatively accurate predictions about the future, in Tetlock's studies, were precisely those who were less certain. The overconfident individuals, who would presumably be more prone to thinking epistemic paternalism would be justified, were the ones with the bad track record. So, when it comes to such complex policy matters, those who would be most disposed toward epistemic paternalism are ironically perhaps the least qualified to do so.

6. Analogs to Old-Fashioned Paternalism and Some Limits

As individuals, we have different tastes, predilections, desires, and values. Mill (1859), p. 66) says in Chapter 3 of *On Liberty*, that "Human nature is not a machine to be built after a model, and set to do exactly the work prescribed for it, but a tree, which requires to grow and develop itself on all sides, according to the tendency of the inward forces which make it a living thing." For him, this makes paternalism unwarranted, as different ways of life are best for different people and, moreover, individuals themselves are (for the most part) in the best epistemic position to determine what is best for them.

The epistemic analog to this is that we all have different sets of total evidence. We have different life experiences, have read different things, talked with different people, and so on. And what matters for rational belief is one's *total evidence*, not a proper subset of it (Kelly, 2008b). Furthermore, pieces of evidence have complex interactions with one another. The more distinct our sets of total evidence are and the more complex the interactions between various pieces of evidence are, the less I can be sure what it is rational for you to believe once a new piece of evidence comes in. This is what goes wrong in the case of Mustard from Section 2: because Mustard is not privy to Peacock's total evidence, his paternalistic intervention has a negative spillover.

[15]This is analogous to a point made by Hayek (1945) in his classic article on the price mechanism. Hayek's point is that the knowledge relevant to constructing a working economic order is essentially distributed across the population. It thus cannot be made available to a group of central planners. While Hayek was specifically focused on economic policy, we might extend this insight more broadly. The evidence relevant to complex questions about policy, in general, is distributed across the population and may not be available to one expert or group of experts. This, in part, explains why prediction markets—which aggregate information from anyone willing to bet—can outperform experts (Surowiecki, 2004).

https://doi.org/10.1017/can.2025.5 Published online by Cambridge University Press

I do not want to overstate the import of this point. Old-fashioned paternalism is often justified, in particular when it comes to children. Thus, it is justifiable to give your 4-year-old broccoli even if he wants to eat candy at the moment. Analogously, in specific cases, it might be appropriate for a mathematics teacher to interfere with her students' inquiry in getting them to come to understand long division. However, old-fashioned paternalism is harder to justify when it comes to adults—and likewise, epistemic paternalism is difficult to justify when it comes to adults, especially in cases where the evidential landscape is complex and the evidence relevant to the proposition in question does not fall neatly within a domain of expertise.

Before concluding, I want to briefly address the case of judges blocking certain pieces of evidence in court, which has been central to defenses of epistemic paternalism (Ahlstrom-Vij, 2013; Goldman, 1991). Judges can make certain pieces of evidence, for instance, about the defendant's purported character or their past crimes, inadmissible in court. Part of the stated rationale for this is that the jury might be unduly prejudiced by such evidence. They might overweight it in their deliberations. If these rules and their justifications are appropriate, then epistemic paternalism seems to be baked into our institutions in an important way.

I want to note two points about these sorts of cases that make them unique. First, when it comes to what kinds of evidence are admissible in court, what is relevant is not only the epistemic position of a particular jury, but also the sorts of downstream incentives that might be created. For example, most obviously, if evidence obtained illegally, without the appropriate warrants, is admissible in court that creates all sorts of perverse incentives for law enforcement and the rest of the population. Similarly, we might worry about evidence of past criminal activity and character traits, and what kinds of perverse downstream incentives would be created if such evidence were admissible in court.

Second, and more importantly, the worry of epistemic spillovers in such cases is of minimal normative importance. What is at hand in a criminal trial is a special context where it is of utmost importance to determine whether the defendant is guilty beyond reasonable doubt. If there are epistemic spillovers—and there might well be—they are normatively insignificant when compared to the matter at hand. If the jury comes to an inaccurate belief about the defendant's character, say, that is much less significant as compared to coming to an accurate decision about guilt. Furthermore, paternalism is restricted to a specific deliberative context—the jury presumably can inquire as they please once the trial is over.

Part of what makes epistemic paternalism about policy matters different is that the spillovers may not be insignificant. And such paternalism, if it is to succeed, needs to be *ongoing*. If the goal is to get the broader public to have an accurate view about some proposition, and policymakers determine that epistemic paternalism is warranted, then their interference with the public's inquiry must be constant. Or, if the target proposition pertains to some "emergency situation," then the interference must last at least till the emergency is over. (This also raises the prospect of the paternalistic policy being revealed after the emergency, which would presumably erode trust in the society's epistemic institutions.)

Epistemic paternalism, then, is a blunt tool for policymakers. It might be useful here to compare blunt medical interventions. Certain drugs or procedures may cure an ailment but cause many negative side effects. Ideally, what medical professionals ought to do is to cure the ailment while minimizing such side effects. If there is a rash on someone's leg, the proper course of action is not amputation but rather some effective ointment. In the epistemic case, we might explore whether there are less blunt tools for expert intervention.

One possibility that has been recently discussed by Begby (2021) is that of evidential preemption. Suppose there is some piece of evidence E that is likely to mislead a novice about whether P because the novice is prone to overweighting E. An expert may preempt this issue by communicating to the novice that E has already been taken into account by the relevant research. So, rather than paternalism, this sort of preemption constitutes *guidance*, and presumably avoids the problems with spillovers I have been discussing here. It also reduces the risk of a loss of trust in expert

https://doi.org/10.1017/can.2025.5 Published online by Cambridge University Press

deliberation and communication. At any rate, if there are less blunt tools of this form that can be used effectively, the arguments presented here suggest that we should do so.

7. Conclusion

Epistemic paternalism may well be justified in specific, highly circumscribed cases. In these cases, the paternalist is in an extremely good epistemic position relative to the target individual, and there are unlikely to be what I have called spillover effects—or such effects are unlikely to be normatively significant. A second-grade mathematics teacher is presumably a good case. Perhaps judges can fall under this bill as well, in specific circumstances. However, when it comes to policy matters of great importance, epistemic paternalism should not be used lightly. That is because, *inter alia*, withholding or suppressing evidence in one area might have unforeseen consequences in another. And subject matter experts cannot rule out such possibilities without risking epistemic trespassing. Thus, I have argued that if experts can find better methods to communicate to the public without compromising on transparency and honesty, they ought to do so. If this article's arguments are plausible, the epistemic burden the paternalist must meet to be justified in departing from such ideals might be higher than it seems at first blush.[16]

Hrishikesh Joshi is an assistant professor of Philosophy at the University of Arizona. His research focuses on social epistemology and philosophy, politics, and economics (PPE).

References

Ahlstrom-Vij, K. (2013). *Epistemic paternalism: A defence*. London: Palgrave Macmillan.

Anderson, E. (2011). Democracy, public policy, and lay assessments of scientific testimony. *Episteme*, 8(2), 144–164.

Applebaum, B. (2019). *The economists' hour*. New York: Little, Brown and Company.

Ballantyne, N. (2015). The significance of unpossessed evidence. *The Philosophical Quarterly*, 65(260), 315–335.

Ballantyne, N. (2019). Epistemic trespassing. *Mind*, 128(510), 367–395.

Ballantyne, N. (2021). The fog of debate. *Social Philosophy and Policy*, 38(2), 91–110.

Begby, E. (2021). Evidential preemption. *Philosophy and Phenomenological Research*, 102(3), 515–530.

Blatz, C. W., & Mercier, B. (2017). False polarization and false moderation: Political opponents overestimate the extremity of each other's ideologies but underestimate each other's certainty. *Social Psychological and Personality Science*, 9(5). https://doi.org/10.1177/1948550617712034.

Braley, A., Lenz, G. S., Adjodah, D., Rahnama, H., & Pentland, A. (2023). Why voters who value democracy participate in democratic backsliding. *Nature Human Behavior*, 7, 1282–1293.

Bullock, E. C. (2016). Mandatory disclosure and medical paternalism. *Ethical Theory and Moral Practice*, 19(2), 409–424.

Bullock, E. C. (2018). Knowing and not-knowing for your own good: The limits of epistemic paternalism. *Journal of Applied Philosophy*, 35(2), 433–447.

Clifford, W. K. (1877). The ethics of belief. *Contemporary Review*, 29, 289–309.

Dasgupta, S. (2022). School in the time of Covid. *Monash Bioethics Review*, 40, 120–144.

Davies, A. (2022). Science communication, cultural cognition, and the pull of epistemic paternalism. *Journal of Applied Philosophy* 40, 65–78. https://doi.org/10.1111/japp.12612.

DiPaolo, J. (2022). What's wrong with epistemic trespassing? *Philosophical Studies*, 179, 223–243.

Goldman, A. (1991). Epistemic paternalism: Communication control in law and society. *Journal of Philosophy*, 88(3), 113–131.

Goldman, A. (2001). Experts: Which ones should you trust? *Philosophy and Phenomenological Research*, 63 (1), 85–110.

Graso, M., Henwood, A., Aquino, K., Dolan, P., & Chen, F. X. (2022). The dark side of belief in Covid-19 scientists and scientific evidence. *Personality and Individual Differences*, 193, 111594. https://doi.org/10.1016/j.paid.2022.111594.

Hardwig, J. (1991). The role of trust in knowledge. *The Journal of Philosophy*, 88(12), 693–708.

Hayek, F. A. (1945). The use of knowledge in society. *The American Economic Review*, 35(4), 519–530.

[16]This article benefited greatly from the detailed comments provided by two anonymous reviewers at the *Canadian Journal of Philosophy*. Thanks also to the editors of this special issue. In developing this article, I have also benefited from discussion with Robin McKenna and feedback from audiences at the PPE Society Meeting at New Orleans, the University of Arizona, the Alabama Philosophical Society, and Arizona State University.

Jackson, E. (2022). What's epistemic about epistemic paternalism? In J. Matheson & K. Lougheed (Eds.), *Epistemic Autonomy*. New York: Routledge.
John, S. (2018). Epistemic trust and the ethics of science communication: Against transparency, openness, sincerity and honesty. *Social Epistemology*, 32(2), 75–87.
Kahan, D. (2012). Why we are poles apart on climate change. *Nature*, 488, 255. https://doi.org/10.1038/488255a.
Kahan, D., & Braman, D. (2006). Cultural cognition and public policy. *Yale Law & Policy Review*, 24(1), 149–172.
Kahan, D., Jenkins-Smith, H., & Braman, D. (2011). Cultural cognition of scientific consensus. *Journal of Risk Research*, 14(2), 147–174.
Kahan, D., Peters, E., Dawson, E. C., & Slovic, P. (2017). Motivated numeracy and enlightened self-government. *Behavioural Public Policy*, 1(1), 54–86.
Kahneman, D. (2011). *Thinking fast and slow*. New York: Farrar, Straus and Giroux.
Kelly, T. (2008a). Disagreement, dogmatism, and belief polarization. *The Journal of Philosophy*, 105(10), 611–633.
Kelly, T. (2008b). Evidence: Fundamental concepts and the phenomenal conception. *Philosophy Compass*, 3(5), 933–955.
Kelly, T. (2014). *Evidence*. Stanford Encyclopedia of Philosophy. Stanford, CA: The Metaphysics Research Lab. https://plato.stanford.edu/entries/evidence/.
Levendusky, M. S., & Malhotra, N. (2016). (Mis)Perceptions of partisan polarization in the American public. *Public Opinion Quarterly*, 80(S1), 378–391.
Marks, J., Copland, E., Loh, E., Sunstein, C., & Sharot, T. (2019). Epistemic spillovers: Learning others' political views reduces the ability to assess and use their expertise in nonpolitical domains. *Cognition*, 188, 74–84.
McKenna, R. (2020). Persuasion and epistemic paternalism. In G. Axtell & A. Bernal (Eds.), *Epistemic paternalism: Conceptions, justifications, and implications*. Washington D.C.: Rowman & Littlefield.
McKenna, R. (2023). *Non-ideal epistemology*. New York: Oxford University Press.
Mercier, H. (2020). *Not born yesterday: The science of who we trust and what we believe*. Princeton, NJ: Princeton University Press.
Mercier, H., & Sperber, D. (2017). *The enigma of reason*. Cambridge, MA: Harvard University Press.
Messina, J. P. (2023). *Private censorship*. New York: Oxford University Press.
Mill, J. S. (1859). On liberty. In J. Gray (Ed.), *On liberty and other essays (2008)*. Oxford, New York: Oxford University Press.
Nguyen, J. (2024). Confidence in Covid-19 models. *Synthese*, 203(4), 1–29.
Page, L. (2022). *Optimally irrational: The good reasons we behave the way we do*. New York: Cambridge University Press.
Pasek, M. H, Ankori-Karlinsky, L.-O., Levy-Vene, A., & Moore-Berg, S. L. (2022). Misperceptions about out-partisans' democratic values may erode democracy. *Scientific Reports*, 12(1), 16284.
Pigou, A. C. (1920). *The economics of welfare*. London: Macmillan.
Sherman, D. K., Nelson, L. D., & Ross, L. D. (2003). Naïve realism and affirmative action: Adversaries are more similar than they think. *Basic and Applied Social Psychology*, 25(4), 275–289.
Surowiecki, J. (2004). *The wisdom of crowds*. New York: Doubleday, Anchor.
Tetlock, P. E. (2005). *Expert political judgment: How good is it? how can we know?* Princeton, NJ: Princeton University Press.
Tetlock, P. E., & Gardner, D. 2015. *Superforecasting: The art and science of prediction*. New York: Crown.
Tsai, G. (2014). Rational persuasion as paternalism. *Philosophy & Public Affairs*, 42 (1), 78–112.
Westfall, J., Van Boven, L., Chambers, J. R., & Judd, C. M. (2015). Perceiving political polarization in the united states: Party identity strength and attitude extremity exacerbate the perceived partisan divide. *Perspectives on Psychological Science*, 10 (2), 145–158. https://doi.org/10.1177/1745691615569849.

Cite this article: Joshi, H. 2025. Science Communication, Paternalism, and Spillovers. *Canadian Journal of Philosophy*: 1–15, doi:10.1017/can.2025.5