

7 What's Epistemic about Epistemic Paternalism?

Elizabeth Jackson

7.1 Introduction

Paternalism is a familiar part of our lives – consider a professor who enforces a no-technology policy for her students, an adult pulled over for not wearing their seatbelt, or a spouse who hides cake from their partner who just started a new diet. Paternalism is the practice of limiting the free choices of agents, without their consent, for the sake of promoting their best interests (see Mill, 1869; Dworkin, 2010; Grill & Hanna, 2018).

There are many strands of paternalism; this chapter focuses on one in particular: *epistemic paternalism*.¹ Roughly, one might think about the distinction this way: regular paternalism aims at (or has the final goal of) improving another's decisions or actions; epistemic paternalism aims at improving another's beliefs. More precisely, epistemic paternalism involves interfering with agents, without their consent, for their own *epistemic* good – e.g. to promote their true beliefs, knowledge, etc.

The aim of this chapter is twofold: (i) to critically examine the concept of epistemic paternalism and (ii) to explore the consequences of normative questions one might ask about it. In Section 7.2, I critically examine several definitions of epistemic paternalism. I argue that many existing definitions are either too broad or too narrow, and I suggest some ways these definitions might be improved. In Section 7.3, I contrast epistemic and general paternalism, and argue that it's difficult to see what makes epistemic paternalism an *epistemic* phenomenon at all. In Section 7.4, I turn to normative questions about epistemic paternalism. I examine different perspectives by which we might evaluate epistemic paternalism, and discuss the literature's assumptions of epistemic consequentialism and veritism. I close in Section 7.5 by comparing and contrasting epistemic paternalism with other phenomena in social epistemology, such as disagreement or testimony. I argue that epistemic paternalism is a uniquely social phenomenon, in a way that e.g. disagreement and testimony are not.

The aim of this chapter is largely clarificatory, rather than an attempt to argue for a single controversial thesis. However, we will address head-on a number of questions at the forefront of the epistemic paternalism

literature that are normally overlooked or quickly brushed aside. As we go, I will flag these questions. Some I will answer or begin to answer, and others I will leave open. Either way, the literature on epistemic paternalism should be paying more attention to them, and eventually address them. Doing so will crucially aid us in answering questions about epistemic paternalism, including questions about whether and how it might be justified.

7.2 Defining Epistemic Paternalism

This section critically examines several definitions of epistemic paternalism in the literature. Goldman (1991) and Lougheed (2021) define epistemic paternalism quite narrowly, focusing on withholding evidence for the sake of promoting true beliefs. Lougheed (2021, p. 261), following Goldman (1991, p. 114), provides the following definition:

If 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, there are instances when Y need not make available to X all of the evidence relevant to Q if doing so will make X more likely to believe the truth about Q.

This suggests the following:

Definition 1: Epistemic paternalism =_{df} (i) withholding evidence from someone, (ii) without their consent, (iii) to make it more likely that they believe truths (or avoid errors).

(i)–(iii), taken as necessary and jointly sufficient conditions for epistemic paternalism, is quite narrow. This isn't necessarily a criticism of Goldman or Lougheed; they have the right to examine and/or attempt to justify a narrower phenomenon. However, insofar as our goal is to capture the class of cases that we would naturally classify as epistemic paternalism, both (i) and (iii) are too narrow.

First, note that condition (iii) seems to assume veritism, the view that believing truth (and avoiding error) is the only final epistemic good. While this is a popular view of epistemic value, it is, nonetheless, controversial (see DePaul, 2001; Carballo, 2018). We can make the definition more ecumenical by revising it in the following way:

Definition 2: Epistemic paternalism =_{df} (i) withholding evidence from someone, (ii) without their consent, (iii) for their own epistemic good.

This definition, unlike the first, is neutral on epistemic axiology – it is consistent with veritism, but also leaves open whether you might engage

in epistemic paternalism to promote goods other than true belief. For example, you might engage in epistemic paternalism to provide another with new, epistemically justified beliefs or knowledge, to increase the level of justification of their already-existing beliefs, or to turn their beliefs into knowledge. I don't see a reason to rule these out as kinds of epistemic paternalism without argument. Veritism is a substantial assumption, and this second definition enables us to avoid assuming it without ruling it out.

I will spend more time on condition (i). Definition 2 claims epistemic paternalism necessarily involves a particular kind of interference: *withholding evidence*. The problem with this condition is that there are a number of cases of epistemic paternalism that aren't a matter of withholding evidence from someone. Consider providing evidence one wouldn't have otherwise had or considered. One might fail to consent to have information that spoils a movie or show, and even express a desire to not have the information (at least at that time, via testimony), but another might give them the information anyway. While undesired, this nonetheless constitutes an epistemic improvement (e.g. more true beliefs/rational beliefs/knowledge) and thus counts as epistemic paternalism. Bullock (2016, 2018) discusses a more serious example: cases of a patient's right not to know about her medical condition if she so chooses. In these cases and others, one may engage in epistemic paternalism by *providing* evidence.

A second kind of epistemic paternalism involves interfering with the way another processes or weighs evidence. For example, when teaching, I might present my class with two philosophical theories, but then strongly emphasize simplicity while purposefully leaving out discussion of the value of explanatory power – making it much more likely they believe the simpler one. I could similarly push skepticism, causing them to value avoiding falsehoods more than they value getting at truths. This could cause them to withhold belief in many propositions, but neglect the epistemic value of true belief. Generally, one can engage in epistemic paternalism by influencing another's *epistemology* – which affects how they process evidence – without changing their first-order evidence.²

Another example of epistemic paternalism involves controlling the order in which someone receives pieces of evidence. Even if the order of evidence wouldn't matter for ideally rational agents, empirical results show that the order in which a normal human reasoner receives different pieces of evidence often affects their conclusion (i.e. the ordering effect).³ For example, when patients are choosing between treatment options, the order in which information is presented may influence their decisions. Bansback et al. (2014) make a case for this and argue that patient decisions can be improved by presenting the most important information first. Thus, one might interfere with inquiry by simply influencing the order in which one receives pieces of evidence.

Other examples of epistemic paternalism include deception and coercive measures, both discussed by Bullock (2018). An example of deception would be teaching a false theory to ultimately facilitate understanding of more complex true theory, such as teaching Newtonian physics before quantum mechanics (2018, p. 434). One might engage in coercive epistemic paternalism by threatening a consequence if someone doesn't form a certain belief. Even if one cannot believe directly based on a threat, they could take action to try to cultivate the belief (2018, p. 435).

The general lesson is this: epistemic paternalism need not merely be about withholding evidence, but can involve a number of other epistemic practices. If we understand 'inquiry' broadly, to include both evidence-gathering and belief-forming practices, then these practices might all be classified as instances of *interfering with another's inquiry*. Given this, we can modify our definition as follows:

Definition 3: Epistemic paternalism =_{df} (i) interfering with someone's inquiry, (ii) without their consent, (iii) for their own epistemic good.

This definition parallels the one found in Ahlstrom-Vij (2013a, p. 51) and Bullock (2018, p. 434). Note that both Ahlstrom-Vij and Bullock have a non-consultation condition rather than a non-consent condition. Bullock (2018, p. 434) explains that an epistemically paternalistic practice "does not consult those interfered with on the issue of whether they should be interfered with in the relevant manner." This is too weak. Suppose I consult with you about some interference I'm considering imposing on you, and you say you don't want me to interfere. I choose to interfere anyway (e.g. maybe I'm threatening you to form a certain belief). This seems like epistemic paternalism, even though I consulted with you. Thus, a non-consent condition seems more plausible than a non-consultation condition.

Thus, definition 3 is about consent, rather than consultation. While I have even utilized definition 3 in my own previous work – Jackson (2020, p. 201) – I now think it may need additional modification. To see why, suppose you decided to write a new book on the philosophy of mind. It is published, and a professor decides to assign it in their philosophy of mind class. As a result of reading the book, students learn a lot of new things about philosophy of mind, and some of them radically change their views.⁴ Now, as a result of writing the book, you've interfered with the inquiry of these students (concerning certain questions in the philosophy of mind). None of these students consented to you writing this book, but (let's suppose) they've epistemically benefitted from reading it. This case seems to fulfill all three conditions, but it's not clear that you have

engaged in epistemic paternalism – writing a new book is just not clearly paternalistic in any way.

This case brings up several issues. First, epistemic paternalism needs to be *intentional*. Note that we are constantly interfering with others' inquiry by, e.g. withholding evidence from almost everyone around us all the time. There are always additional ways that we could epistemically improve others. But it doesn't seem like we are constantly engaging in epistemic paternalism. Nonetheless, as noted above, many cases of epistemic paternalism seem to involve intentionally withholding important evidence from someone. Part of what explains the difference is that accidental or unintentional paternalism seems impossible. Thus, we might edit our definition to reflect that epistemic paternalism is *intentional*.⁵ Thus, withholding evidence only counts as epistemic paternalism if it is done intentionally.

This doesn't fully deal with our case, though – you likely wrote the book to epistemically improve *someone*, even if not those students in particular. It doesn't seem that paternalism, even if intentional, needs to be directed at a particular person. A government might build a wall around a dangerous part of a mountain hike to prevent people from exploring that area. The wall isn't built to keep out certain individuals in particular, but still seems paternalistic.

One might instead think that, in the philosophy of mind case, the lack-of-consent condition isn't met. Maybe the students do, in fact, consent by, e.g. enrolling in the relevant course. Suppose instead a student finds the book randomly in the library and starts reading it. While the student didn't consent to you writing the book, they make the decision to pick it up and read it, so this doesn't count as epistemic paternalism. This suggests that consenting to *the book's being written* isn't the relevant object of consent. Generally, we should pay more attention to the consent condition. What counts as consent? And *what* exactly do the agents in question need to consent to?

More attention also needs to be paid to what counts as an *interference*.⁶ Note that interfering with another's inquiry comes in degrees. On the more innocent end of the spectrum, there is sharing a fact with someone that they didn't ask about – consider the talkative neighbor who always wants to tell you about their day, their job, or their kids. Then there's handing someone a book or pamphlet, or recommending they read something. Then there are more serious interferences, like epistemic manipulation: sharing certain facts but purposefully leaving others out, or intentionally presenting certain arguments to your class so they come to have a particular view. On the most extreme end of the spectrum, there's e.g. serious acts of lying and brainwashing, and epistemic threats and coercion. We can suppose that, in all these cases, the consent condition isn't met and there

is an epistemic improvement, so they are all at least candidates for epistemic paternalism.

But it's not clear that we should count all these cases as epistemic paternalism. If your neighbor is a talker, you may learn more than you want to know from them, but is this really paternalistic? Is handing your friend an interesting book to read paternalistic? Maybe the interference has to be sufficiently *significant* to count as epistemic paternalism. This suggests:

Definition 4: Epistemic paternalism =_{df} (i) intentionally and significantly interfering with someone's inquiry, (ii) without their consent, (iii) for their own epistemic good.

I worry this definition is still not sufficiently precise. (What counts as a "significant interference"? What counts as consent?) This points to the following first open question:

Open Question 1: when defining epistemic paternalism, how should we understand *consent* and *interference*?

And the general discussion in this section to a second, related question:

Open Question 2: what is the best way to define epistemic paternalism, so that it is neither too broad nor too narrow?

In my view, something close to definition 4 is promising and perhaps can deal with the tricky cases by being clearer about the consent and interference conditions. Nonetheless, this discussion shows that even basic questions about the definition of epistemic paternalism need more attention and refinement.

7.3 Epistemic and General Paternalism

This section continues our exploration of the nature of epistemic paternalism. More specifically, it contrasts epistemic paternalism with general paternalism and assesses what makes the former an *epistemic* phenomenon at all. While epistemic paternalism is generally taken to be a type of general paternalism, it is also assumed to be a unique subset that is distinguishable from cases of non-epistemic paternalism. This section argues that either the distinction is much more slippery than it appears *prima facie*, or cases of epistemic paternalism are extremely rare.

Consider again our definition: the interference must happen *for* the inquirer's good. As a preliminary note, it is natural to read this as claiming that the inquirer's epistemic good is a *reason* for the interference. What kind of reason exactly? Metaethicists distinguish different kinds of

reasons – three of which are *motivating* reasons (facts for which someone phi-s), *normative* or justifying reasons (that count in favor of phi-ing), *explanatory* reasons (that explain why someone phi-ed). Since our goal here is simply to define epistemic paternalism, I think the best candidate here is motivating reasons – the epistemic benefit is the fact for which they interfered, and normative reasons for epistemic paternalism come in when we examine the question of whether epistemic paternalism is justified (e.g. in the next section).⁷

One of the main things that distinguishes epistemic and general paternalism is that the interference must be done for the inquirer's *epistemic* good. In other words, one of the primary factors that characterizes epistemic paternalism is that it is motivated in a distinctly epistemic way. This raises the question: need the interferer be motivated *only* by the inquirer's epistemic good? What if they act *partially* for the inquirer's epistemic good? Many cases in the actual epistemic paternalism literature are arguably instances of the latter. For example, consider a commonly used case: a judge withholding evidence from a jury to raise the probability they will come to the right verdict (maybe the judge has good evidence that the jury will weigh the evidence improperly). Presumably, the judge might do this in part for the epistemic good involved (the jury's getting a true belief, justified belief, knowledge, etc.) but in real-life cases, a major part of the judge's motivation is moral: to convict the guilty and to let the innocent go free. This latter thing is only contingently connected to the jury's having true beliefs, and if (for some odd reason) the judge thought that the jury's having false beliefs would lead to convicting the guilty and the innocent going free, the judge would likely not be motivated to interfere in the same way. This suggests that the epistemic is not the judge's primary motive.

Goldman (1991), who pioneered the philosophical interest in epistemic paternalism, is sensitive to this point. After stating the definition of epistemic paternalism, he says, "the restriction to the epistemic viewpoint is again important. In legal settings, for example, there are many non-epistemic reasons for refusing to provide relevant evidence to jurors" (114). However, in many of the examples of epistemic paternalism he provides, it is not at all clear that the agents involved are motivated purely by epistemic factors. For instance, he considers epistemic paternalism in education curriculum, noting that, in health classes, we don't give "equal time to drug pushers" (121), which has an obvious non-epistemic motivation (keeping kids from doing drugs). He also mentions the battle over whether creationism should be taught in schools, noting that excluding creationism from curriculum might be a case of epistemic paternalism. However, he even admits that, in this case, non-epistemic constitutional issues about including religion in public education muddy the waters (122). He also mentions epistemic paternalism

in commercial advertising, to combat false or deceptive advertising. While he claims that the goal in this is “to keep [buyers] from believing untruths about commercial products” (122), this is clearly not the only goal, and most people care about this goal because they care about the actual products people will eventually buy. Similar considerations apply to most of Goldman’s other examples, including interfering with what news items are covered by the media, and how they are covered – people’s beliefs from watching and reading the news have significant non-epistemic effects. In fact, Ahlstrom-Vij (2013a, p. 117, 134) invokes an alignment condition for this very reason: on his view, one of two jointly sufficient conditions for justified epistemic paternalism is that the epistemic and non-epistemic reasons must be normatively aligned (e.g. both permit the interference). One reason he needs this condition is because of how frequently epistemic paternalism is tied up with, and affects, the non-epistemic.

This raises the question: are there *any* cases of “pure” epistemic paternalism, where someone is motivated to interfere with another’s inquiry solely on epistemic grounds? These cases turn out to be quite rare.⁸ A potential set of cases involves abstract topics in science, philosophy, etc. that have little practical import. For example, Bullock (2018, p. 434) discusses an instructor acting paternalistically to teach her students quantum mechanics. Something similar may happen in some philosophy classes – e.g. a professor engages in epistemic paternalism to help her students better understand the realist/nominalist debate, and does so purely to facilitate their understanding of the issues, rather than for any practical or moral reasons. These seem like plausible cases where one might be motivated to interfere with another’s inquiry purely for the sake of promoting true/justified beliefs or knowledge, and not for downstream non-epistemic effects.

Note also that not all paternalism regarding abstract concepts is purely epistemic, since we might sometimes be motivated to help someone understand something abstract for a practical or moral reason (e.g. to help them get into grad school or to get a good teaching evaluation). Of course, this point applies in the opposite direction to the cases above as well; an interference might have non-epistemic benefits, but our motives for interfering could still be solely epistemic.⁹ Nonetheless, I suspect this normally isn’t the case; again, we care about a jury’s beliefs because we want to free the innocent and convict the guilty; we care about our children’s beliefs about drugs so they don’t do them, etc.

The general lesson is that, if epistemic paternalism requires acting *only* for another’s epistemic good, it is probably a very rare phenomenon, even if we understand epistemic value quite broadly (as including true beliefs, justified beliefs, knowledge, understanding, etc.). This isn’t great news for those working on epistemic paternalism because it means that

(i) real-life cases of epistemic paternalism are virtually non-existent, and
 (ii) many of the literature's supposed examples of epistemic paternalism, such as the cases of juries in courtrooms, are not actually cases of epistemic paternalism.

Alternatively, one might maintain that an act can count as epistemic paternalism if it is done *partially* for another's epistemic good, but practical and moral factors can also be part of the motivation. There are several issues with defining epistemic paternalism this way, however. First, this move strays from current literature. Goldman (1991), Pritchard (2013), and Bullock (2018) are all sensitive to the idea that paternalism motivated partially by the epistemic and partially by other, non-epistemic factors doesn't seem distinctly epistemic.¹⁰ Of course, one could push back on this, but would be out of step with current literature.

Second, if we allow epistemic paternalism to be motivated by both epistemic and non-epistemic reasons, then it also becomes less clear what makes epistemic paternalism distinctly epistemic. Yes, it involves interfering with inquiry, but are those interferences of any special epistemic interest if they are done partially or even mostly for practical or moral reasons? For example, if I interfere with a juror's evidence partially to help them come to know who did it, but also out of moral concern from the guilty to be punished and the innocent to be set free, it's not clear that my interference ought to be classified as pure epistemic paternalism (it might be classified as a pseudo-epistemic paternalism). You also might wonder if the weight of each reason matters: what if the epistemic is only a small part of the reason I interfere, and the moral is my primary motive? Maybe an act counts as epistemic paternalism if the epistemic reason is weightier than the other ones. Or maybe it counts if the epistemic reason is sufficient, on its own, to motivate the interferer. These possibilities are underexplored.

Finally, in many real-life cases, the epistemic motivation is (only or primarily) valuable *instrumentally*. Consider, for instance, why we care about buyer's beliefs about products they might purchase, children's beliefs about the effects of drugs, or a jury's beliefs about who committed a crime. It is because of the downstream moral and practical effects of these beliefs; we rarely purely care about the beliefs themselves. Thus, in many realistic cases, it's not clear that the distinctly epistemic factors are central to the motivation for the interference *at all*. Again, recall our example in which a jury's having false beliefs would lead to convicting the guilty and the innocent going free. While the judge may not *lie* to the jury (note that she also has moral and practical reasons not to lie), she may withhold evidence even if it makes the jury epistemically worse off, if it leads to practically and morally good consequences.

To sum up, when defining epistemic paternalism, we are faced with a dilemma. If epistemic paternalism requires acting only for another's epistemic good, then we're spending a lot of intellectual energy and ink

analyzing a practice that rarely occurs (and most of our examples don't actually apply). However, if an action can count epistemic paternalism if it is motivated by both epistemic and non-epistemic reasons, then it becomes less clear what distinguishes epistemic paternalism from regular paternalism, why it would deserve its own literature and analysis, and why its justification would differ from that of regular paternalism.¹¹ This leaves us with at least two open questions:

Open Question 3: Does epistemic paternalism require acting *only* for another's epistemic good?

Open Question 4: What, if anything, distinguishes epistemic paternalism from regular paternalism?

Now, we turn to normative evaluations of both paternalism and epistemic paternalism.

7.4 Evaluating Epistemic Paternalism

We've now discussed the nature of epistemic paternalism and ways that it contrasts with general paternalism. While many questions are unanswered, and I've provided reason to doubt that the distinction is well-defined, I'll proceed as if we have some way to differentiate the two. More specifically, in this section, I will (modestly) assume that, if an action is primarily motivated by non-epistemic factors, we ought not count it as epistemic paternalism (even if this means that epistemic paternalism is a rare phenomenon). Now, we turn to normative questions.

As a preliminary note, it is instructive to examine justifications for regular paternalism and epistemic paternalism side by side, because lessons and concepts from the longstanding, mature paternalism literature can be applied to questions about epistemic paternalism, which is newer and less developed. For example, in cases of general paternalism, there is widely taken to be a presumption of non-interference. We should err on the side of not interfering with the free choices of others, unless we have an overriding reason to do so (Mill, 1869, Dworkin, 2010, sec. 3). The "burden of proof," then, is on the person engaging in paternalism to justify their interference. One might wonder: does this presumption hold in the epistemic case? Interestingly, most of the epistemic paternalism literature thus far involves arguments that it is justified, and almost no one has argued that epistemic paternalism is always or almost always impermissible (Bullock, 2018 being the notable exception). Since so many people seem to be in favor of epistemic paternalism (apparently epistemic libertarians are quite rare!), this might suggest that the presumption of non-interference doesn't hold in the epistemic case. At the same time, philosophers like to argue for controversial theses, so this observation

about the literature may not mean much. It is nonetheless worth exploring whether the presumption that applies to general paternalism also applies to epistemic paternalism.¹²

Now, when it comes to justifications of general and epistemic paternalism, there are at least four questions one might ask:

- Q1. Is paternalism all-things-considered justified?
- Q2. Is paternalism epistemically justified?
- Q3. Is epistemic paternalism all-things-considered justified?
- Q4. Is epistemic paternalism epistemically justified?

Here, I understand all-things-considered justification to include moral, practical, epistemic reasons, and it may also include other types of value, like aesthetic reasons. (Some, e.g. Feldman, 2000, are skeptical of all-things-considered justification because they think different types of reasons aren't commensurable.) Note that there are questions in addition to Q1–Q4 that isolate each of these types of value, e.g. is epistemic paternalism morally (or practically) justified? Note also that, assuming epistemic paternalism is a subset of general paternalism, then a negative answer to Q1 entails a negative answer to Q3, and a negative answer to Q2 entails a negative answer to Q4. Similarly, a positive answer to Q4 entails a positive answer to Q2, and a positive answer to Q3 entails a positive answer to Q1. Finally, it is crucial to be clear about which question is of interest, rather than simply speaking of epistemic paternalism's "justification." It could very well be that epistemic paternalism is, e.g. epistemically justified but morally unjustified, or practically justified but not all-things-considered justified. For example, Bullock (2018)'s argument that epistemic paternalism is *morally* unjustified might not give us reason to doubt Goldman's (1991) argument that it is *epistemically* justified.

Several of those who have written on the justification of epistemic paternalism thus far have focused on Q3.¹³ However, this strikes me as a debate that isn't especially fruitful, because it seems difficult, if not impossible, to justify epistemic paternalism on all-things-considered grounds. This is because, as Pritchard (2013) and Bullock (2018) point out, whatever epistemic gain supposedly motivates epistemic paternalism – e.g. true belief, justified belief, knowledge – will almost always be outweighed by moral considerations in favor of personal autonomy and sovereignty. (Recall here that, if we understand epistemic paternalism as primarily motivated by epistemic considerations, we cannot count the moral or practical value conferred by the content of the beliefs; it is purely the epistemic value of the true/justified beliefs/knowledge). Bullock (2018, pp. 442–443) gives the following illustrative example:

Suppose, for example, that I play a series of physics lectures to you whilst you are sleeping, with the intention that you subconsciously

learn quantum mechanics. I have good reason to think this will be effective. You happen to have no interest in quantum mechanics and the facts that you learn have no bearing on your wellbeing. Is this interference justified on balance? ... It seems intuitive ... that the loss to personal sovereignty is in fact a weightier concern than the gain in knowledge: indeed, it looks as though you would be morally correct to admonish me for my secretive interferences even if you wanted to learn about quantum mechanics.

I agree with Bullock's assessment of this case – it is both morally and all-things-considered impermissible for me to sneak the headphones on you while you are sleeping, simply because I want you to learn about quantum mechanics. It is hard to see how the epistemic gain of doing so could outweigh the moral losses. This leads to our next open question:

Open Question 5: Can epistemic paternalism ever be morally or all-things-considered justified?

In general, it is hard to see how epistemic paternalism could be morally or all-things-considered justified, especially if it requires acting solely for another's epistemic good. Part of this depends on what counts as a significant interference, per definition 4. For instance, if I give you a book about quantum mechanics instead, it's much less clear that this is morally and all-things-considered impermissible.¹⁴ However, Bullock might reply that this interference isn't sufficiently significant, so it doesn't count as epistemic paternalism.

But there is one notable exception to this that Bullock's argument overlooks – the case of *nudges*. (This is another example of how epistemic paternalism can fruitfully borrow from regular paternalism.) The concept of a nudge was pioneered by Sunstein and Thaler (2003). They argue that, in certain situations, we are forced to present choices to others in a certain way. In these cases, no matter what we do, we will have some influence on their decisions. Sunstein and Thaler argue that, for this reason, we might as well present the choices in a way that makes it more likely that people will choose what is best for them. This is called a nudge. For example, in a cafeteria, students are more likely to choose foods that are at eye level. One might “nudge” these students by putting healthier foods at eye level, making it more likely that they will pick the healthy option. *Something* must be put at eye level, so, the reasoning goes, we might as well choose the healthier food. Sunstein and Thaler argue that nudges are a case of paternalism that even libertarians should be happy with – it is difficult to see why nudges would be unjustified. This raises the question: is there such a thing as *epistemic nudges*? For example, when presenting people with a body of information, we have to present some information first. As discussed in Section 7.2, humans are subject to ordering

effects, so that the order in which pieces of evidence are presented affects people's ultimate judgments. We might present certain information first to facilitate our audience's understanding, even though this might change what they ultimately conclude. After all, we have to present it in *some* order, so we might as well do what is epistemically best for our audience. Nudges may be cases of uncontroversially morally/all-things-considered justified epistemic paternalism. Generally, epistemic nudges strike me as a fruitful area for further research.¹⁵ Further, they strike me as one of the rare cases where there is a moral or all-things-considered justification for epistemic paternalism.

Because morally and all-things-considered justified epistemic paternalism is plausibly rare, a potentially more fruitful debate involves Q4, which is the original focus of Goldman's (1991) paper: is epistemic paternalism *epistemically* justified? This question is about the epistemic justification of a particular practice, namely, interfering with inquiry. This assumes that epistemic norms can guide behavior. Further, 'epistemic justification' (or 'epistemically justified') is used here in a non-standard way – it doesn't pick out the thing that turns true unGettiered belief into knowledge.¹⁶ Here, 'justification' indicates when a practice, on balance, promotes epistemic goods.

This sheds light on why many in the epistemic paternalism literature have either implicitly or explicitly adopted a version of epistemic consequentialism.¹⁷ Epistemic consequentialism is the view that epistemic goods are more fundamental than epistemic obligations – what you epistemically ought to do is promote certain epistemic goods. Many candidate deontological epistemic norms – e.g. believe in accord with your evidence, believe truths, don't believe contradictions, have probabilistic credences, etc. – are not norms that guide action, but are evaluative or teleological norms that apply to belief. It is unclear what deontological epistemic principles would govern action; at least, these principles are not frequently discussed by epistemologists. However, it nonetheless seems epistemically good to take actions that promote valuable epistemic states, such as true belief, justified belief, coherent belief, etc. Thus, one way you might epistemically justify a particular practice is by arguing that it will likely result in many epistemically good states and/or enable the avoidance of epistemically bad states, or, more simply, that it maximizes expected epistemic value (Greaves, 2013).

A potential paper explores what the deontological, epistemic, action-guiding norms might be (if they exist), and how those affect the normative status of epistemic paternalism. Many in this literature simply assume a version of epistemic consequentialism. However, note that we need not be epistemic consequentialists across the board to make sense of Q4. We need only assume epistemic consequentialism about epistemic norms for *action*; we can be deontologists about epistemic norms for *belief*. We can still maintain that deontological norms such as: *believe in*

accord with your evidence or *have probabilistic credences* apply directly to attitudes themselves. However, when it comes to how we should *act*, epistemically, consequentialist norms kick in: we should act in ways that, e.g. maximize expected epistemic value.

Nonetheless, this still leaves the following question:

Open Question 6: Can epistemic norms guide action?

If the answer is no, then we cannot evaluate epistemic paternalism from an epistemic point of view at all. And several philosophers, including Feldman (2000), Kelly (2002, fn. 30), Berker (2018), and Simion (2018), argue that there aren't epistemic reasons for action. This leads to another potential dilemma, this time concerning the normative evaluation of epistemic paternalism. Either we evaluate epistemic paternalism from a moral or all-things-considered perspective, in which case it is hard to see how epistemic paternalism is justified (this is essentially Bullock's (2018) argument, nudges being the key exception), or, we evaluate it epistemically, in which case we must controversially assume that that epistemic norms can guide action. The latter might not be so bad, insofar as it is reasonable to think some epistemic norms guide certain kinds of behaviors, such as how we get evidence (e.g. inquiry, evidence gathering) and what we do with our evidence (e.g. critical reasoning, reflection on our evidence).¹⁸ But with either natural way of evaluating epistemic paternalism, we are left with surprising or controversial results. Now, we turn to ways that epistemic paternalism compares and contrasts with other questions in social epistemology.

7.5 Epistemic Paternalism and Social Epistemology

Goldman (1991) mentions several times that epistemic paternalism falls under social epistemology, or what he calls *social epistemics*. He explains, "Social epistemics studies the veritistic properties of social practices, or institutional rules that directly or indirectly govern communication and doxastic decision" (120). Goldman and O'Connor (2019), in the *Stanford Encyclopedia of Philosophy* entry on 'Social Epistemology', define social epistemology as "an enterprise concerned with how people can best pursue the truth (whichever truth is in question) *with the help of*, or *in the face of*, others." Later in that same article, they categorize some of the central topics in social epistemology as testimony, disagreement, how we should identify and respond to expert belief, and epistemic injustice.

This suggests a noteworthy sense in which epistemic paternalism differs from other questions in social epistemology. Many of the topics Goldman and O'Connor mention – disagreement, testimony, epistemic injustice, and responding to expert opinion – can be framed as (at least primarily) concerning the question, "What should *I* believe?" The disagreement

literature asks what I should believe in the face of peers who disagree with me. The epistemology of testimony is about whether, and under what conditions, I should believe a testifier. Questions about experts also concern the question of what I should believe in response to experts (or so-called experts). Even epistemic injustice can be seen as largely concerning the question of what I should believe – in response to the testimony of marginalized groups.

Interestingly, epistemic paternalism notably differs from these other topics, and is, in some sense, more social. This is because epistemic paternalism concerns our *epistemic obligations to others* – not merely what we should believe when we encounter other epistemic agents. It is about how we ought to affect – or not affect – the beliefs of others. This complex issue doesn't merely boil down to the question of what a single individual should believe, but involves tricky considerations about how agents should treat each other, epistemically. One reason this is notable is because epistemic obligations to other people are rarely discussed, in any literature; in fact, many assume that our obligations to others are merely moral ones.¹⁹ Further, it is noteworthy that many issues in social epistemology are less social than one might have thought, and can be grouped alongside issues in traditional epistemology, at least insofar as they all concern the question, "What should I believe?" Of course, the fact that epistemic paternalism goes beyond this question might explain some why some of the issues mentioned above might crop up – e.g. whether there are epistemic reasons for action, whether epistemic consequentialism is true, and how we should understand epistemic justification when it comes to our obligations to others. Nonetheless, I think this shows that we can further divide social epistemology into distinct and interesting categories, and some traditional issues in social epistemology might be notably "less social" than others. This brings us to a final open question:

Open Question 7: What is the best way to define social epistemology, and how should we categorize its topics?

While, unlike the other questions, this question is not necessarily crucial for those writing on epistemic paternalism to address, I nonetheless think it alludes to stimulating and significant general issues in social epistemology.

7.6 Conclusion

My primary goal has been to critically examine the concept of epistemic paternalism and to survey and evaluate various normative questions we might ask about it. While this paper has raised a lot of problems for both the concept of epistemic paternalism and its evaluation, I nonetheless hope

that I've enabled clearer thinking about epistemic paternalism. This includes what epistemic paternalism is, what controversial commitments it may carry (e.g. potentially that there are epistemic reasons for action, and maybe certain strands of epistemic consequentialism), and what commonly made assumptions in the literature are unnecessary (e.g. veritism). Ultimately, I hope this chapter facilitates the aim of accurately answering questions about what kinds of epistemic paternalism (if any) are justified, under what circumstances they are justified, and in what sense they are justified.

Acknowledgments

Thanks to Kirk Lougheed and Jon Matheson for helpful comments on an earlier draft. Thanks to Seth Lazar, Justin D'Ambrosio, Nic Southwood, Matthew Kopec, Klaas Kray, Chris Dragos, and audiences at the 2019 Canadian Philosophical Association, Australian National University, and Michigan State University for valuable discussion and feedback. Research on this chapter was supported by Australian Research Council Grant D170101394.

Notes

- 1 Discussions of epistemic paternalism include Goldman (1991), Ahlstrom-Vij (2013a, 2013b), Pritchard (2013), Ridder (2013), Bullock (2018), Croce (2018). See also Bernal and Axtell (2020) for an edited volume on epistemic paternalism.
- 2 See Jackson (forthcoming) and Jackson and Turnbull (forthcoming) for a further discussion of the ways one's broader epistemic situation can affect one's beliefs without affecting one's evidence.
- 3 See, for example, Walker et al. (1972), Dean (1980), Hogarth and Hillel (1992), Wiegmann et al. (2012).
- 4 Thanks to Frank Jackson for suggesting this case to me.
- 5 Thanks to Jon Matheson.
- 6 Thanks to Kirk Lougheed.
- 7 Thanks to Seth Lazar for helpful discussion; see Ahlstrom-Vij (2013a, p. 113).
- 8 Thanks to Seth Lazar for helpful discussion.
- 9 Thanks to John Matheson.
- 10 Ahlstrom-Vij (2013a, pp. 117, 134) argues that for epistemic paternalism to be *justified*, epistemic and non-epistemic reasons need to be aligned in a particular way. But his point isn't that non-epistemic reasons are frequently part of the *motivating* reason for epistemic paternalism; his point is about *normative* reasons (117).
- 11 One possibility is that both general and epistemic paternalism essentially concern the same phenomenon, but are simply evaluated from different points of view; the former evaluates it morally and practically, and the latter evaluates it epistemically. (Thus, we are only concerned with Q1 and Q2 from the next section.) While I'll proceed as if epistemic paternalism is a unique phenomenon, this suggestion warrants further exploration. Thanks to Kirk Lougheed and Jon Matheson for helpful discussion.

- 12 Thanks to Nic Southwood.
- 13 Including Pritchard (2013) and Ahlstrom-Vij (2013a). Bullock (2018) focuses on moral justification for epistemic paternalism (see her fn. 7), but my comments in this paragraph also apply to her view.
- 14 Thanks to Jon Matheson.
- 15 Thanks to Matt Kopec. On epistemic nudges, see Meehan (2020).
- 16 Thanks to Pamela Robinson.
- 17 For recent defenses of epistemic consequentialism, see Singer (2018a, 2018b). See also Ahlstrom-Vij and Dunn (2018).
- 18 See Tidman (1996), Hookway (1999), Friedman (2019).
- 19 One exception is Basu (2019), who discusses a potential epistemic obligation to others: not to wrong others in what we believe about them. (Note also that Basu advocates for moral encroachment, so this obligation is both epistemic and moral, since, according to moral encroachment, the moral can affect epistemic rationality). Another potential exception in the philosophy of testimony literature is the obligation that speakers have to hearers. However, this may not be distinctly epistemic either (e.g. lying or intentionally misleading another is often taken to be *morally* impermissible, not merely epistemically impermissible). For more on testimony and social obligations, see Goldman (1999) and Lackey (2008). Thanks to Jon Matheson.

References

- Ahlstrom-Vij, K. (2013a). *Epistemic paternalism: A defence*. London: Palgrave Macmillan.
- Ahlstrom-Vij, K. (2013b). Why we cannot rely on ourselves for epistemic improvement. *Philosophical Issues*, 23, 276–296.
- Ahlstrom-Vij, K., & Dunn, J. (Eds.). (2018). *Epistemic consequentialism*. Oxford: OUP.
- Bansback, N. L., Li, L. C., Lynd, L., & Bryan, S. (2014). Exploiting order effects to improve the quality of decisions. *Patient Education and Counseling*, 96(2), 197–203.
- Basu, R. (2019). What we epistemically owe to each other. *Philosophical Studies*, 176, 915–931.
- Berker, S. (2018). A combinatorial argument against practical reasons for belief. *Analytic Philosophy*, 59(4), 427–470.
- Bernal, A., & Axtell, G. (2020). *Epistemic paternalism reconsidered: Conceptions, justifications, and implications*. Lanham, MD: Rowman & Littlefield.
- 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.
- Carballo, A. P. (2018). Good questions. In K. Ahlstrom-Vij & J. Dunn (Eds.), *Epistemic consequentialism*. Oxford: OUP.
- Croce, M. (2018). Epistemic paternalism and the service conception of epistemic authority. *Metaphilosophy*, 49(3), 305–327.
- Dean, M. L. (1980). “Presentation order effects in product taste tests.” *The Journal of Psychology*, 105(1): 107–110.

- DePaul, M. (2001). Value monism in epistemology. In M. Steup (Ed.), *Knowledge, truth, and duty: Essays on epistemic justification, responsibility, and virtue* (pp. 170–183). Oxford: Oxford University Press.
- Dworkin, G. (2010). Paternalism. In E. N. Zalta (Ed.), *The Stanford encyclopedia of philosophy*. <https://plato.stanford.edu/entries/paternalism/>.
- Friedman, J. (2019). “Inquiry and belief.” *Noûs* 53(2), 296–315.
- Goldman, A. I. (1991). Epistemic paternalism: Communication control in law and society. *The Journal of Philosophy*, 88(3), 113–131.
- Goldman, A. I. (1999). *Knowledge in a social world*. Oxford: Oxford University Press.
- Goldman, A. I., & O’Connor, C. (2019). Social epistemology. In E. N. Zalta (Ed.), *The Stanford encyclopedia of philosophy*. <https://plato.stanford.edu/entries/epistemology-social/>.
- Greaves, H. (2013). Epistemic decision theory. *Mind*, 122(488), 915–952.
- Grill, K., & Hanna, J. (2018). *The Routledge handbook of the philosophy of paternalism*. New York: Routledge.
- Hogarth, R. M. E., & Hillel, J. (1992). Order effects in belief updating: The belief-adjustment model. *Cognitive Psychology*, 24(1), 1–55.
- Hookway, C. (1999). Epistemic norms and theoretical deliberation. *Ratio*, 12(4), 380–397.
- Feldman, R. (2000). “The ethics of belief.” *Philosophy and Phenomenological Research* 60(3): 667–695.
- Jackson, E. (2020). Epistemic paternalism, epistemic permissivism, and stand-point epistemology. In A. Bernal & G. Axtell (Eds.), *Epistemic paternalism reconsidered: Conceptions, justifications, and implications* (pp. 201–215). Lanham, MD: Rowman & Littlefield.
- Jackson, E. (Forthcoming). A defense of intrapersonal belief permissivism. *Episteme*.
- Jackson, E., & Turnbull, M. G. (Forthcoming). Permissivism, underdetermination, and evidence. In C. Littlejohn & M. Lasonen-Aarnio (Eds.), *The Routledge handbook of the philosophy of evidence*. New York: Routledge.
- Kelly, T. (2002). The rationality of belief and some other propositional attitudes. *Philosophical Studies*, 110(2), 163–196.
- Lackey, J. (2008). *Learning from words: Testimony as a source of knowledge*. Oxford: Oxford University Press.
- Lougheed, K. (2021). Epistemic paternalism, open group inquiry and religious knowledge. *Res Philosophica*, 98(2), 261–281.
- Meehan, D. (2020). Epistemic vices and epistemic nudging: A solution? In A. Bernal & G. Axtell (Eds.), *Epistemic paternalism reconsidered: Conceptions, justifications, and implications* (pp. 247–261). Lanham, MD: Rowman & Littlefield.
- Mill, J. S. (1869). *On liberty*. London: Longman, Roberts & Green.
- Pritchard, D. (2013). Epistemic paternalism and epistemic value. *Philosophical Inquiries*, 1(2), 9–37.
- Ridder, J. D. (2013). Is there epistemic justification for secrecy in science? *Episteme*, 10(2), 101–116.
- Simion, M. (2018). No epistemic norm for action. *American Philosophical Quarterly*, 55(3), 231–238.

- Singer, D. J. (2018a). How to be an epistemic consequentialist. *Philosophical Quarterly*, 68(272), 580–602.
- Singer, D. J. (2018b). Permissible epistemic trade-offs. *Australasian Journal of Philosophy*, 97(2), 281–293.
- Sunstein, C. R., & Thaler, R. H. (2003). Libertarian paternalism is not an oxymoron. *University of Chicago Law Review*, 70(4), 1159–1202.
- Tidman, P. (1996). Critical reflection: An alleged epistemic duty. *Analysis*, 56(4), 268–276.
- Walker, L., Thibaut, J., & Andreoli, V. (1972). Order of presentation at trial. *Yale Law Journal*, 82, 216–226.
- Wiegmann, A., Okan, Y., & Nagel, J. (2012). Order effects in moral judgment. *Philosophical Psychology*, 25(6), 813–836.