

Epistemic paternalism and epistemic value*

Duncan Pritchard

Abstract: My concern is with the merits (or otherwise) of a form of paternalism which is specifically *epistemic*, one that mirrors familiar forms of paternalism in the ethical and political spheres. To this end, I critically explore Alvin Goldman's seminal discussion of this topic, which runs along veritistic lines. While broadly sympathetic to the general position that Goldman develops in this regard, I argue that in order to properly evaluate the potential of epistemic paternalism we need to broaden our focus away from veritism and consider also the special value that attaches to certain epistemic standings.

1. *Paternalism, epistemic paternalism and veritism*

The legitimacy, or otherwise, of paternalistic practices is widely discussed in the ethical and political literature. Very roughly – we do not need to get into the details for our purposes paternalistic practices are practices which limit the liberty of agents in order to promote their best interests.¹ As the name suggests, the root idea behind paternalism is the kind of supervisory control that a parent exercises over their child, but the main focus of philosophical interest in paternalism is concerned with its institutional application. Examples abound of such practices, particularly when it comes to legal restrictions which are imposed by the state.

Take the law that one should wear seat belts. For the sake of argument, let us stipulate that the sole reason for such a law is to protect the individual welfare of the drivers and passengers concerned. (As it happens, this is probably false, in that those not wearing a seat belt can in practice be a danger to others particularly other passengers in the vehicle

* Invited Paper.

¹ For a helpful overview of the recent literature on paternalism, see Dworkin (2010). For some of the key texts relevant to this debate, see Kleinig (1983), Sartorius (1983), Feinberg (1986), and Van DeVeer (1986).

and not just themselves).² Such a law imposes a limit on one's freedom, in that one must abide by the law regardless of whether one consents to it (and if one doesn't abide by it then one's freedoms will be further limited), but the point of the law is to promote the welfare of the very drivers and passengers whose freedom is being limited.

A range of issues are raised by paternalistic practices. Which bodies have the authority to undertake such practices? How does one weigh-up the loss of liberty against the goods which accrue to the agent concerned in virtue of the practices? While it is in the nature of such practices that the agent concerned does not give her explicit consent, to what extent should these practices be such that the agent ought to be in some sort of in principle position to offer consent (for example, such that the following counterfactual is true: had she been consulted, provided with the relevant facts, and fully rational, then she would have consented)? And so on.

My interest in this paper is not, however, with the general issues raised by paternalism, but rather in exploring the particular question of the legitimacy, or otherwise, of paternalistic practices when it comes to specifically epistemic goods – *viz.*, *epistemic paternalism*. Such a question is central to the topic of the epistemology of liberal democracy, since it raises the possibility that two key goods of a liberal society – *viz.*, individual liberty and autonomy – might be curtailed in the service of specifically epistemic goals. But while there has been much discussion of paternalism in moral and political debate, there has been very little exploration of paternalism in epistemology. Indeed, I do not think it would be an exaggeration to say that Alvin Goldman's seminal 1991 paper, 'Epistemic Paternalism: Communication Control in Law and Society', is really the only substantive discussion of this issue in the epistemological literature.³

As Goldman makes clear in this paper, it is undeniably the case that epistemic paternalism is explicitly practiced in several quarters. One very clear example which he explores is that of the role of evidence in criminal

² In the UK at least, part of the rationale for requiring back-seat car passengers to wear seat belts is the potential harm a back-seat passenger without a seat belt could cause in an accident to others in front of them in the vehicle.

³ Kristoffer Ahlstrom is preparing a monograph on epistemic paternalism, and when this appears this will of course substantially further the debate. I am grateful to Kristoffer for passing onto me a draft manuscript of one of his papers on this topic (Ahlstrom 2011). In addition, one of my PhD students Shane Ryan is working on a thesis project which explores what he refers to as *epistemic environmentalism*, which is a view that is closely related to epistemic paternalism. Ryan envisages a campaigning and, ultimately, policing movement which mirrors, in the epistemic sphere, the political doctrine of environmentalism. So construed, epistemic environmentalism would be a particularly radical form of epistemic paternalism.

trials within the US legal system. He points out that the rationale offered for a number of the rules which make up the Federal Rules of Evidence are explicitly epistemic, and in some cases involves a degree of epistemic paternalism. For example, these rules sometimes forbid, on epistemic grounds, jurors from gaining access to some types of evidence, on the grounds that such evidence would mislead them and thereby frustrate the epistemic ends of the criminal justice process. Here is Goldman:

[...] it is apparent that the framers of the rules, and judges themselves, often wish to protect jurors in their search for truth. If, in the framers' opinion, jurors are likely to be misled by a certain category of evidence, they are sometimes prepared to require or allow such evidence to be kept from the jurors. This is an example of what I shall call *epistemic paternalism*. The general idea is that the indicated rules of evidence are designed to protect jurors from their own "folly," just as parents might keep dangerous toys or other articles away from children, or might not expose them to certain facts (Goldman 1991: 118).

Note (as Goldman himself points out), the primary recipients of the epistemic goods (putatively) being generated by these paternalistic practices are not the jurors themselves but rather the defendants, and hence this example is in this regard at least somewhat different to standard examples of paternalism in the literature (like the 'seat belt' example above). Even so, this case shares enough key features with standard cases of paternalism to qualify as a *bona fide* example of epistemic paternalism. In particular, in this case we have an established practice of intervening with regard to the information available to jurors, without the consent of the parties concerned, in order to promote certain epistemic goods, the promotion of which would otherwise be undermined. Of course, even if epistemic paternalism were not in fact practiced, we could still ask the question of whether it should be practiced and, if so, what form it should take. But that it is in fact practiced gives the question of whether epistemic paternalism is ever legitimate an 'applied' edge, since it means that what stance we take on this question can have ramifications for current practice.

In evaluating epistemic paternalism, it is of course vital that we have some conception of the epistemic good(s) in play relative to which we undertake the target evaluation. On this score, what little discussion there has been of epistemic paternalism in the literature has tended to take it as given that the epistemic good should be understood along *veritistic* lines. According to veritism, ultimately all that matters from a purely epistemic point of view is just true belief (and the avoidance of false belief), with all

other epistemic properties being merely instrumentally valuable relative to the good of true belief.⁴ The foremost exponent of veritism is Goldman, and so it is no surprise that he applies this view to the issue of epistemic paternalism.⁵ For example, he writes that in evaluating epistemic paternalism we need to “inquire into the general circumstances in which” it “has good or bad epistemic consequences, i.e., good *veritistic* outcomes.” (Goldman 1991, 124) Others have followed him on this score. In his recent defence of epistemic paternalism, Kristoffer Ahlstrom (2011, 4) explicitly evaluates such paternalism relative to “paradigm epistemic goals”, where these are in turn characterized as “the formation of true belief and the avoidance of false belief.”

In this paper, I will be arguing that we should evaluate epistemic paternalism relative to a broader conception of the epistemic good than that which is supplied by veritism. In order to see why this is necessary, we first need to consider some key threads in the recent debate about epistemic value.

2. *Epistemic value*

There has been a huge growth in interest in questions about epistemic value in the recent literature, to the extent that some have spoken of epistemology’s ‘value-turn’.⁶ One defect in quite a lot of this literature is a failure to notice a crucial ambiguity in the very idea of epistemic value. Construed one way, to say that something (e.g., a belief) is of epistemic value is to attribute a particular kind of value to it (i.e., a kind of value which is epistemic). Construed another way, to say that something is of epistemic value is to say that it is something which is epistemic (e.g., an

⁴ Henceforth, for the sake of simplicity, and since nothing of consequence hangs upon it, I will simply take veritism to be the claim that the ultimate and fundamental epistemic good is true belief, and drop the *caveat* about veritism being in addition concerned with the avoidance of false belief.

⁵ Goldman defends veritism in a number of places, but see especially Goldman (1999; 2002) and Goldman & Olsson (2009). Elsewhere see, e.g., Pritchard, Millar & Haddock (2010, ch. 1) and Pritchard (2011) I’ve referred to veritism as *epistemic value T-monism*, to make clear that this is a type of epistemic value monism where there is only one fundamental epistemic good, and where the fundamental epistemic good in question is true belief rather than some other epistemic standing. For more on veritism, see David (2001), and the exchange between David (2005) and Kvanvig (2005).

⁶ See Riggs (2008). For two useful surveys of recent work on epistemic value, see Pritchard (2007) and Pritchard & Turri (2011). See also Pritchard, Millar & Haddock (2010, ch. 1).

epistemic standing, like knowledge) and which is also valuable.⁷ Very roughly, one way of putting this point is in terms of *epistemic value* on the one hand, and the *value of the epistemic* on the other. It is important to keep these two notions apart since neither entails the other.

On the one hand, that something is specifically of epistemic value does not entail that it is valuable along some other axis of evaluation – such as aesthetically – still less does it entail that it is valuable *simpliciter* (we will be considering what being valuable *simpliciter* might involve below). Perhaps, for example, true belief is epistemically valuable, but not at all valuable in any other respect, and moreover not valuable *simpliciter*. On the other hand, to say that something epistemic is valuable without qualification does not entail that its value is purely of the epistemic kind. Perhaps, for example, knowledge is valuable in ways that aren't specifically epistemic. It could be, for instance, that it has ethical value; indeed, it might well be valuable *simpliciter*.⁸

In talking about epistemic value we must thus be careful to be clear exactly what it is that we have in mind. A good illustration of why this is important can be found by considering the so-called 'swamping' problem that has been much explored in the contemporary literature.⁹ According to this problem, certain epistemological proposals are unable to account for the putative intuition that knowledge is of greater value than mere true belief, and hence are to that extent problematic.

Reliabilism is often cited as a key case in point in this regard.¹⁰ For, so the argument goes, why should the fact that a true belief is reliably formed add any value to that true belief? After all, since we only care about reliable belief-forming processes because we care about having true beliefs, why should it matter whether a belief is reliably formed so

⁷ The distinction set out here is related to the point that Geach (1956) makes about 'predicative' and 'attributive' adjectives. In particular, from 'x is a big flea' it does not follow that 'x is a flea' and 'x is big', since the claim being made is only the attributive claim that x is big *for a flea*. In contrast, it does follow from the predicative claim that 'x is a red flea' that 'x is a flea' and 'x is red'. In Geach's terms, then, my claim is that 'x is epistemically valuable' is ambiguous between an attributive and a predicative reading.

⁸ I explore this distinction between epistemic value and the value of the epistemic in more detail in Pritchard, Millar & Haddock (2010, ch. 1). See also Pritchard (2011). Although, as noted above, a lot of epistemologists fail to mark this distinction, some are clearly sensitive to it. See, for example, Sosa (2007, ch. 4).

⁹ For some key discussions of this problem, see Jones (1997), Swinburne (1999), Riggs (2002a; 2002b), Kvanvig (2003) and Zagzebski (2003). See also Pritchard (2007a; 2007b; 2011) and Pritchard, Millar & Haddock (2010, ch. 1).

¹⁰ See, for example, Zagzebski (2003).

long as it is true? An analogy that is often offered to illustrate this point is the ‘coffee cup’ case (e.g., Zagzebski 2003). We care about reliable coffee-making machines because we care about good coffee. But if we are presented with two identical cups of coffee, only one of which came from a reliable coffee-making machine, we will surely value both cups the same, even despite their very different sources. What goes for coffee from a reliable coffee-making machine is held to also apply to true belief that is produced by a reliable belief-forming process.

The more general point being made here is not of course specific to reliabilism. For what underlies the swamping problem is the idea, roughly, that if we only value something as a means to something else, then the presence of the former can contribute no additional value when the latter is also present. In the epistemic case this means that if we only value a certain epistemic standing as a means to a further, more fundamental, epistemic standing, then the presence of the former can contribute no additional value when the latter is also present. So, for example, if justification is only valuable as a means to true belief, then that one’s true belief is in addition justified does not make it any more valuable than a corresponding mere true belief.

So construed, veritism is clearly the real target of the swamping problem, since this precisely does hold that what is valuable from an epistemic point of view is ultimately down to the propagation of true belief. In cases where one already has a true belief, then, there will be a standing problem of explaining why any additional epistemic standing (such as being reliably-formed) should, by veritistic lights, contribute any further value. Goldman, as the chief proponent of veritism, has developed (with Erik Olsson) some sophisticated ways of dealing with the swamping problem.¹¹ As I explain at length elsewhere see Pritchard (*forthcomingd*) I don’t think these responses to the swamping problem work. Crucially, however, I maintain that with the distinction just noted between the two different ways of thinking about epistemic value in play, Goldman doesn’t need to undertake such theoretical manoeuvres in order to evade the swamping problem, in that a very straightforward response is available to him.

In order to see this, consider first the challenge that Goldman thinks the swamping problem poses for his view. He writes that the swamping problem

¹¹ See Goldman & Olsson (2009; cf. Olsson 2007; 2009). For a critical discussion of this paper, see Pritchard (*forthcomingd*).

[...] can be used to test the adequacy of accounts of knowledge. Suppose an analysis of knowledge is incompatible with knowledge having an added value [*relative to mere true belief*]. That would be a pretty strong argument against the adequacy of that analysis. (Goldman & Olsson 2009: 22)

The claim in play here is ambiguous though. The thesis that one's theory of knowledge should be at least compatible with knowledge being of more value than mere true belief sounds right, but note that the intuitiveness of this claim rests on our treating the additional value in question as being more than just epistemic. For consider what the corresponding thesis would be if we focused on specifically epistemic value *viz.*, that one's theory of knowledge should be at least compatible with knowledge being of more specifically *epistemic* value than mere true belief. Is *that* claim intuitive? I don't see why it should be. In particular, it seems like a moot theoretical issue whether knowledge adds specifically epistemic value to mere true belief. In particular, it seems that a claim of this sort is just what a veritist ought to be suspicious about, since for them what is of ultimate and fundamental epistemic value is just true belief.

Crucially, however, the swamping problem is targeted not at the issue of whether knowledge is *in general* more valuable than mere true belief, but whether it is specifically *epistemically* more valuable. Indeed, if this were not so, then the argument offered in support of the swamping problem would be invalid. Consider again how the swamping problem is meant to work, using the specific target of veritism to illustrate the point. First, we have a general claim about value, to the effect that if the value of a property possessed by an item (e.g., being produced by a reliable coffee-making machine, or being the product of a reliable belief-forming process) is only instrumental value relative to a further good (e.g., good coffee, or true belief) and that good is already present in that item, then this property can confer no additional value on that item. Next, we have the veritist thesis that, roughly, all that really matters from an epistemic point of view is true belief. This gets re-cast, plausibly, as the claim that all other epistemic standings that a belief might bear over and above being true are only of instrumental epistemic value relative to the fundamental and ultimate good of true belief. Putting these two claims together we are meant to get the conclusion that knowledge, as an epistemic standing that a belief might bear over and above being true, is of no more value than mere true belief.

We can express this argument more formally as follows:

The Swamping Problem for Veritism

- (SP1) If the value of a property possessed by an item is only instrumental value relative to a further good and that good is already present in that item, then this property can confer no additional value to that item.
- (SP2) Any epistemic standing that a belief might bear over and above being true is only of instrumental epistemic value relative to the good of true belief. [Veritism]
- (SC) Knowledge is no more valuable than mere true belief. [From (SP1), (SP2)]

The conclusion of this argument clearly is in conflict with the thesis that we saw Goldman endorsing above, to the effect that one's epistemology ought to be at least consistent with knowledge being of more value than mere true belief. Thus cast, then, the swamping problem does seem to pose a serious challenge for veritism.

Crucially, however, note that as that as it stands (SC) does not follow from (SP1) and (SP2). (P2), after all, is clearly concerned with a specifically epistemic kind of value, since that is the very nature of the veritism thesis. It follows that the most that can be concluded from (SP1) and (SP2) is the strikingly weaker (SC*):

- (SC*) Knowledge is no more *epistemically* valuable than mere true belief. [From (SP1), (SP2)]

But far from being a claim which is problematic for veritism, (SC*) seems to merely be a direct consequence of the view, and hence nothing for the veritist to be concerned about. In particular, (SC*) is entirely in keeping with the thought that we saw Goldman articulating above to the effect that one's theory of knowledge ought to be compatible with knowledge being of more value than mere true belief. This is because (SC*) does not rule-out the (very plausible) possibility that knowledge is more valuable than mere true belief in respects which aren't specifically epistemic, perhaps by being ethically valuable or by being of greater practical utility.¹²

¹² I think this point highlights a key difference between the swamping problem and the so-called 'Meno problem'. While both are concerned with the greater value of knowledge over mere true belief, the latter is clearly a more general problem since it is not targeted at the greater specifically *epistemic* value of knowledge. That the *Meno* problem is focussed on the more general value of knowledge relative to mere true belief is clear from the kind of diagnostic line that Socrates considers in response to it, such that we should take into account the greater *practical* value that knowledge has over mere true belief in virtue of the fact that it is 'tethered' to the facts. I explore this point more fully in Pritchard, Millar & Haddock (2010, ch. 1) and Pritchard (2011).

The relevance of this point for epistemic paternalism is that even if we grant, with Goldman and others, the veritist claim that it is true belief which is of ultimate and fundamental epistemic value, there is still a further question to be asked about what the value is more generally of the epistemic standings our beliefs possess. Take knowledge, for example, and suppose it is true as some have argued – that knowledge has a prudential value that mere true belief lacks.¹³ It would then follow that insofar as one is persuaded by the thought that epistemic paternalism is sometimes legitimate, then there is a *prima facie* motivation to consider favouring practices which promote knowledge and not just mere true belief, even if it also true that from a purely epistemological point of view knowledge is no more valuable than mere true belief.

So once we understand that veritism is a claim about specifically epistemic value, then it follows that it is at least an open question whether there are epistemic standings, over and above mere true belief, which the epistemic paternalist should be concerned to promote. There is a further – and related – reason why even the veritist should be willing to explore the wider value of epistemic standings other than mere true belief. As noted above, that something is of specifically epistemic value does not entail that it is valuable along other axes of evaluation, much less does it entail that it is valuable *simpliciter*. So far I've treated this idea of 'value *simpliciter*' as intuitive and so not offered a further characterisation, but now is a good juncture to say more about it.

In principle at least, we can identify for any domain some axis of evaluation relative to which we assess items that fall within that domain. So, for example, we could imagine a domain of evaluation which applies to some wicked activity like serial killing in which we assess agents as good *qua* serial killers (e.g., for their ingenuity and cunning, *à la* Hannibal Lecter), or as bad *qua* serial killers (e.g., those serial killers which are clumsy and ineffective). Or we could imagine a domain of evaluation which applies to some rather trivial activity like train spotting in which we assess agents as good *qua* train spotters (e.g., great memory for trains, lots of commitment to train spotting, good use of relevant observational skills, etc.), or bad *qua* train spotters (e.g., poor memory for trains, very little commitment to train spotting, poor use of relevant observational skills, etc.).

¹³ This is roughly the claim considered by Socrates in the *Meno*, when he talks of how the fact that knowledge is true belief which is 'tethered' to the facts means that it has a greater practical value over mere true belief. See also Williamson (2000, ch. 1) and Pritchard, Millar & Haddock (2010, ch. 5).

We noted above that it doesn't follow from the fact that something is good, *qua* epistemic, that it is good in any other way, much less that it is good *simpliciter*. We can now get a better sense of what is at issue here. For while there might be some domains of evaluation in which what is good relative to that domain is good *simpliciter*, clearly not all domains of evaluation are of this kind. That some agent or act is good relative to the domain of serial killing or train spotting does not intuitively make it at all good *simpliciter*. In contrast, what is good in other domains of evaluation does seem to be good *simpliciter*. The clearest case of this is the ethical good. If a certain action is good from an ethical point of view, then that surely means that to some extent at least the action is good *simpliciter*.

Note that this is not to say that ethically good actions are thereby all-things-considered good *simpliciter*, since that is a further issue which depends on the kinds of things one allows as being good *simpliciter*. So, for example, suppose one held that only what is ethically good or aesthetically good is good *simpliciter*. There would then be a potential trade-off, when it came to all-things-considered value, between the ethical good and the aesthetic good in terms of the ultimate evaluation of the item in question (with, perhaps, the ethical and the aesthetic having different relative weightings in this trade-off). One couldn't thus infer from the fact that something is ethically good that it is all-things-considered good. Nonetheless, the ethical good is at least the kind of thing that enters into the equation when it comes to determining the all-things-considered good, unlike the kinds of goodness that are specific to particular domains of evaluation (as just illustrated by the 'serial killing' and 'train spotting' domains of evaluation).

Is the epistemic domain of evaluation like the ethical domain of evaluation or more like, say, the train spotting domain of evaluation? That is, is the fact that something is of epistemic value the kind of thing that goes into the mix when one determines whether the epistemically valuable item is of all-things-considered value? In order to answer this question, we need to plug-in a specific account of epistemic value, so let's consider the issue with veritism in play. Our question now becomes whether, for instance, an action which promotes true belief, and which is thus epistemically good by veritistic lights, is in the market to be good *simpliciter*.

One can see how a case could be made for such a claim, at least once we remember that the goodness being attributed is defeasible (and possibly not *pro tanto*). Suppose, for example, that one is presented with an action which, it is stipulated, promotes true belief (and hence which is epistemically good by veritistic lights), but which doesn't promote any

other good (e.g., it isn't ethically good, or prudentially good, or aesthetically good, and so on). Isn't the epistemic goodness of this action at least some (defeasible) reason to think the action good *simpliciter*? In contrast, if all one knows about an action is that it promotes the 'goods' of serial killing (and does not promote any other good), then this is surely no reason at all for thinking the action good *simpliciter* (indeed, it is a reason to think the action *isn't* good *simpliciter*). If that's right, then the epistemic good does have more in common with the ethical (or, say, the aesthetic) good, than it does the train spotting or serial killing good.

No doubt many would be unpersuaded by an argument of this sort, but I think one can at least see that there is a *prima facie* plausibility to this line of reasoning. If epistemic paternalism is to be motivated purely in the case of true belief, however, then I think it is important that a case is made by the veritist for this thesis, since otherwise it will be unclear why an agent's liberty should be constrained on an epistemic basis. Liberty, after all, is surely an ethical good, and hence something which is plausibly good *simpliciter*. So if true belief is not good *simpliciter*, why should it be cited as a basis on which liberty might be constrained?

But even if one is persuaded by the idea that what is epistemically good is good *simpliciter*, it still remains that even for the veritist it is an open question whether the goodness of epistemic standings is exhausted by the epistemic good. We considered above the possibility that knowledge has a prudential value that mere true belief lacks. The prudential good might well be good *simpliciter*, but suppose that the additional (non-epistemic) value of knowledge over mere true belief is not (merely) prudential but ethical (as some have claimed; we will be looking at one key argument for this thesis in a moment). If any species of good is a candidate to be good *simpliciter*, it is surely the ethical good, as we noted above. But that would mean that knowledge is good *simpliciter*, in the sense set out above such that its goodness goes into the mix when we determine the all-things-considered value of the item in question. Knowledge would then be an epistemic standing that we would have reason to promote, perhaps more reason to promote, over and above mere true belief. And yet the reason why we would want to promote it over mere true belief would not be because of its additional epistemic value (by veritist lights it doesn't have any additional epistemic value), but because it is valuable in a way that mere true belief is not (i.e., ethically valuable, where this is also valuable *simpliciter*).

The key point is that there is nothing in veritism which goes against a claim like this. That is, as things stand there are a range of theoretical

options on the table, depending on how one thinks of epistemic goodness and how one thinks of the goodness of the epistemic. One question is whether epistemic goodness is goodness *simpliciter*. A second question is whether there is a goodness of epistemic standings which is more than just epistemic goodness. And a third question is whether there is a goodness of epistemic standings which is more than just epistemic goodness and which is goodness *simpliciter*. Depending on how one answers these three questions, one will have a very different view about epistemic paternalism.

For example, if one thinks that the goodness of the epistemic is just the epistemic good, but that the epistemic good is not good *simpliciter*, then, as noted above, one will have a problem even motivating epistemic paternalism in the first place. For given that liberty and autonomy are arguably good *simpliciter*, why should they be limited in the service of purely epistemic ends if the epistemically good is not good *simpliciter*? But even if one does think that the epistemic good is good *simpliciter*, it still remains that there may be epistemic standings which are of greater value than is captured by the epistemic good (i.e., because they are also good in other, non-epistemic, ways), and if so epistemic paternalism will need to consider such standings as well as (or perhaps instead of) the standings which constitute the fundamental epistemic good.

In what follows we will set to one side the question of whether the epistemic good as the veritist conceives of it is good *simpliciter*. Instead, our issue will be whether there are epistemic standings whose goodness is not purely epistemic and where this goodness is – as, arguably, it is in the ethical case – goodness *simpliciter*. As we will see, a strong case can be made that understanding fits the bill in this regard. But if that's right, then in evaluating a strategy of epistemic paternalism we ought to not be merely focussed, as veritism contends, on the propagation of true belief, but also (perhaps instead, though I won't be pressing this matter here) on the propagation of higher epistemic standings like understanding.

3. *The relative value of true belief, knowledge and understanding*

My immediate concern, then, is to argue that some epistemic standings, over and above mere true belief, have a goodness which is good *simpliciter*. One way of arguing for this claim is via the thesis that knowledge, unlike mere true belief (or indeed any epistemic standing which falls short of knowledge) is ethically good.

A number of authors have advanced this claim, but to cut to the quick let us give our own presentation of it:

The Ethical Goodness of Knowledge

(KP1) Achievements are successes that are because of one's ability.

(KP2) Knowledge is cognitive success that is because of one's cognitive ability.

(KC1) Knowledge is cognitive achievement. [From (KP1), (KP2)]

(KP3) Achievements are ethically good.

(KC2) Knowledge is ethically good. [From (KC1), (KP3)]

Let's take a moment to look at the motivation for each of these claims.

The motivation for (KP2) comes from robust virtue-theoretic accounts of knowledge.¹⁴ According to virtue-theoretic theories, knowledge is the product of cognitive ability, or virtue.¹⁵ The mere conjunction of cognitive success (i.e., true belief) and cognitive ability will not suffice for knowledge, however, since it is very easy to 'Gettierize' such a conjunction by imagining a case where the cognitive success is only luckily related to the cognitive ability. This is why robust virtue epistemologists demand that the cognitive success should be *because of* the cognitive ability in play.

With knowledge so construed, however, it is natural to regard it as merely falling under a general type of performance where one's success is creditable to (i.e., because of) one's ability and not other factors (such as luck). This more general type of performance is achievement, where knowledge is now just a specifically cognitive type of achievement. We thus get the argument from (KP1) and (KP2) to (KC1).

Now (KP3) might look to be the hard sell here, in that it is far from obvious that achievements are ethically good. But on reflection this does seem to have a fair degree of plausibility. Ethics, broadly conceived anyway (i.e., such that it is more than just morality), is concerned with the good life, the life of flourishing. Here is the question. Are achievements, so conceived, a key part of such a life? Well, plausibly they are. For imagine a life full of success and ability, but where the successes were never creditable to the agent's ability but instead due to, say, divine intervention, or dumb luck. Wouldn't such a life be impoverished as a result?

¹⁴ For some of the key defences of robust virtue epistemology, see Sosa (1988; 1991; 2007; 2009a), Zagzebski (1996; 1999) and Greco (2003; 2007; 2008; 2009a; 2009b; 2009c).

¹⁵ The identification of cognitive ability with intellectual virtue is thought by some in the literature to be contentious, the thought being that there is much more to intellectual virtue than mere cognitive ability. Since nothing hangs on this point for our purposes, I will be setting this potential complication to one side in what follows.

More generally, all other things being equal, wouldn't one prefer to succeed because of one's ability rather than because of luck or the interference of others? I think the intuition in play here is that achievements reflect our agency genuinely – rather than merely apparently – acting on the world, where this is something that is a key ingredient to a good life.¹⁶

With (KC1) and (KP3) in play, however, then the ultimate conclusion of this argument follows, such that knowledge is of ethical value. It is clearly implicit here that ethical goodness – or at least the kind of ethical goodness at issue in the argument – is goodness *simpliciter*, in the sense specified above that it is a consideration that goes into the mix when we consider the all-things-considered goodness of something. (Indeed, I take it that the claim in play here is that this goodness *simpliciter* is a *pro tanto* goodness, but I won't press this point here).

Moreover, although the argument as expressed doesn't make this explicit, it is also clear that this way of thinking about the value of knowledge is meant to explain why knowledge is distinctively valuable in a way that that which falls short of knowledge isn't.¹⁷ For example, mere cognitive success, like the mere conjunction of cognitive success and cognitive ability, lacks the distinctive properties that mark out cognitive achievement (and thus, on this view, knowledge), and which make the latter uniquely (at least compared with lesser epistemic standings anyway) deserving of this particular kind of goodness. No wonder, then, that we as epistemologists have been so concerned to focus on knowledge to the exclusion of other epistemic standings.

Unfortunately, I don't think this argument for the ethical value of knowledge works, though ultimately it isn't that important to what I'm up to in this paper whether or not this argument is successful. I will explain why this is the case below, but for now let me briefly review what I think is wrong and what I think is right about this argument.

¹⁶ For more on this point, see Pritchard (2010).

¹⁷ That is, it answers a challenge which I have elsewhere termed the 'tertiary' value problem for knowledge. The 'primary' value problem for knowledge is the *Meno*-style challenge of explaining how knowledge is more valuable than mere true belief. (As noted in endnote 12, one should not equate the primary value problem with the swamping problem – though many do – since only the latter is specifically concerned with the greater *epistemic* value of knowledge over mere true belief). The 'secondary' value problem for knowledge is the challenge to explain how knowledge is more valuable than that which falls short of knowledge. Finally, the 'tertiary' value problem is the challenge to explain why knowledge is more valuable, as a matter of kind and not merely degree, than that which falls short of knowledge. For further discussion of this taxonomy of value problems for knowledge, see Pritchard (2009*b*) and Pritchard, Millar & Haddock (2010, ch. 1). See also Pritchard (2007) and Pritchard & Turri (2011).

Let's start with what seems to be the weakest point – *viz.*, (KP3). I think that as achievements are characterised in (KP1), this claim is not very plausible. For note that all kinds of activities will satisfy the rubric for achievements laid down by (KP1). In particular, some very trivial activities, such as merely raising one's arm, will count as achievements by the lights of (KP1). But I think this is to unduly stretch the ordinary language notion of an achievement. In any case, there surely is no intuition that merely raising one's arm is something that is of ethical value, and hence good *simpliciter*.

But that problem needn't mean that the underlying claim at issue in (KP3) is without merit. Perhaps it just needs qualification. For example, consider *strong achievements*, where these are successes that are because of one's ability and which either (i) involve a significant level of skill, or (ii) involve the overcoming of a significant obstacle to one's success. Merely raising one's hand is not a strong achievement, though raising one's hand when one is recovering from an operation on one's arm might be. Conversely, that a great tennis player like Novak Djokovic can strike a fantastic shot with ease does not disqualify it as a strong achievement, in virtue of the level of skill being exhibited.¹⁸

So suppose we replace (P3) with (P3*):

(P3*) Strong achievements are ethically good.

That is, suppose we argue that it is only strong achievements which have the goodness at issue in the argument above. I suggest that (P3*) is at least more compelling than (P3), in that the kind of performance that someone who is recovering from an operation on their arm undertakes when lifting their arm is more plausibly part of the life of flourishing than the kind of performance undertaken by someone who, in normal circumstances, simply raises their arm.

Replacing (P3) with (P3*) will not help the proponent of this argument though, in that with this premise so formulated it's no longer clear that the rest of the argument follows. After all, even if one is persuaded by the general claim that knowledge is a form of achievement, in the sense of a cognitive success that is because of one's cognitive ability, one will surely

¹⁸ One's concern with (KP3) might not just with the idea of easy achievements though, since one might also be concerned that this claim is in tension with the idea of *wicked* achievements, like the 'achievements' of the serial killer. I explore this concern in some detail in Pritchard, Millar & Haddock (2010, ch. 2), and ultimately claim that it is a much less pressing objection to (KP3) than the objection just considered regarding trivial (and unskilful) 'achievements'. See also Pritchard (2010).

find the claim that knowledge is a *strong* cognitive achievement hard to swallow. After all, such a conception of knowledge seems unduly restrictive given how essentially passive and artless much of our acquisition of knowledge (e.g., of our immediate environment) seems to be.

The proponent of this argument for the ethical goodness of knowledge is thus faced with a dilemma. On the one hand, insofar as they have an argument which motivates the idea that cognitive achievements are ethically valuable, then they are restricted to a conception of cognitive achievements on which it isn't plausible that knowledge is a cognitive achievement, so conceived. Or else they have a plausible story about how knowledge should be understood as a cognitive achievement, but don't have a compelling account available to them of why knowledge, so conceived, should be thought to be ethically valuable.

Rather than trying to resolve this dilemma – which I think is irresolvable, though that's a matter for another occasion¹⁹ – I think we should instead attend to an interesting issue that surfaces out of this dilemma. For while the idea that knowledge, *qua* cognitive achievement, is ethically valuable has been seen to be contentious, the idea that strong cognitive achievements in virtue of being a type of strong achievement *simpliciter* are ethically valuable is still in play. And, in fact, that claim is rather plausible.

Take first the claim that strong achievements *simpliciter* are ethically valuable. Imagine, for example, a life filled with successes but where those successes are down to luck rather than one's ability. Or imagine a life which, while full of achievement, only ever involves the weak kind of achievement which involves neither a significant level of ability nor the overcoming of a significant obstacle. Wouldn't such lives be substantially impoverished as a result?²⁰

Now, of course, the special value being assigned to strong achievements *simpliciter* here is defeasible, in that other factors could ensure that all-things-considered any particular strong achievement is not valuable (or of limited value, or even disvaluable). (That said, it may be that the special value of strong achievements, while defeasible, is nonetheless *pro tanto*, though I won't press this point here). But this claim is enough for our purposes, for the point remains that strong achievement is the kind of thing which is ethically good, and thus good *simpliciter*.

The issue then rests on whether the special value that attaches to

¹⁹ I discuss this dilemma in more detail in Pritchard (2009b; 2012) and Pritchard, Millar & Haddock (2010, ch. 2).

²⁰ For more on this point, see Pritchard (2010).

strong achievements *simpliciter* carries over to the specific category of strong cognitive achievements. There is every reason to think that it does. Indeed, a life of flourishing surely involves agents employing not just moral virtue but also intellectual virtue. So insofar as strong achievements *simpliciter* are a key part of such a good life, then one would expect this to include strong cognitive achievements also, where these are creditable to the agent's exercise of intellectual virtue.

There is of course more to be said on this point, but I hope I have done enough here to at least make a *prima facie* case for the claim that strong cognitive achievements have a special ethical value. One way of illustrating the importance of this point for our purposes is to consider the special value of understanding why such-and-such is the case (henceforth just 'understanding').²¹ As I've argued elsewhere see Pritchard (2009a) and Pritchard, Millar & Haddock (2010, ch. 4) understanding, unlike knowledge, essentially involves strong cognitive achievement.

For example, while knowledge can be gained passively and/or artlessly, understanding always requires either overcoming a significant obstacle or else the exercise of a significant level of cognitive skill. A child may passively come to know that his house has burned down because someone authoritative person told him (like a parent), but he doesn't thereby understand why it burnt down. Indeed, a child may passively come to know via the same route that the house burnt down because (say) of faulty wiring, but unless he is able to get an intellectual grip on how faulty wiring can cause a fire, then he won't thereby have any understanding of this event.²² Conversely, while understanding can sometimes come to one immediately, without one having to do anything specific to gain it, this is because one is bringing to bear special cognitive skills. Sherlock Holmes, for example, might well understand immediately why the house is alight, but that's just because of his prodigious cognitive skills.²³

If understanding does involve strong cognitive achievement, and if

²¹ What I say here about understanding why such-and-such is the case can be adapted to apply to other forms of understanding, such as understanding a subject matter, but this raises complications that I have not the space to cover here, which is why I am focusing on this specific type of understanding.

²² I have argued elsewhere that one can also have understanding while lacking the corresponding knowledge, on account of how understanding, unlike knowledge, is compatible with a certain kind of epistemic luck. See Pritchard (2009a) and Pritchard, Millar & Haddock (2010, ch. 4).

²³ This is not to suggest that only someone with prodigious cognitive skills can gain immediate understanding. The point is just that a significant level of cognitive skill is required.

one grants the earlier point that strong cognitive achievements have a special ethical value, then this can go some way to explaining the commonly held conviction amongst epistemologists that understanding is of special value.²⁴ Of particular interest for our purposes, however, is the idea that if understanding has this special value, then it ought to have a role to play in an evaluation of any instance of epistemic paternalism.

4. *Epistemic paternalism, reconsidered*

Let's review where we are. We have seen that in talking of epistemic value we need to be careful to distinguish between, on the one hand, the value of the epistemic and, on the other hand, a specific kind of value which is epistemic. Veritism is specifically a thesis about the nature of epistemic value, and we saw that this provides the view with a straightforward way of dealing with the swamping problem (on account of how the greater value of knowledge over mere true belief does not entail that knowledge is of greater epistemic value than mere true belief).

Once we mark this distinction, however, then this raises a number of questions which veritism must confront if it is to motivate a form of epistemic paternalism. One immediate question concerns the veritistic motivation for epistemic paternalism. If the veritist wants to motivate a strategy of epistemic paternalism purely on the basis that this promotes the epistemic good as the veritist conceives of it, then it is vital that the veritist can show that what is epistemically good is good *simpliciter*. For without that thesis there is no obvious reason why an agent's liberty should be constrained in order to further such epistemic ends. Fortunately for the veritist, we saw above that there is at least a *prima facie* argument for this thesis.

The more fundamental question that is raised by this distinction, however, is that even if the epistemically good is good *simpliciter*, it still remains that the value of certain epistemic standings may not be a specifically epistemic value. Moreover, the value that attaches to these epistemic standings might be weightier than that which attaches to mere true belief. If that is right, then the veritistic focus when it comes to epistemic paternalism on whether such a strategy promotes true belief may be unduly restrictive, in that in evaluating a strategy of epistemic paternalism

²⁴ See, for example, Elgin (1996; 2004; 2009), Zagzebski (2001), Kvanvig (2003) and Riggs (2009).

we should also be considering the further issue of whether this strategy promotes these other epistemic standings.

This point was illustrated above by considering the putative ethical value of certain epistemic standings. We saw that while the case that has been made for the ethical value of knowledge in the recent literature is not ultimately persuasive, much of the argumentative support for this claim can be adapted to make a case for the ethical value of understanding. If that's correct, then in evaluating the merits of a strategy of epistemic paternalism we should consider not only whether it promotes true belief, but also whether it promotes understanding too (indeed, we might well think we should be more concerned about whether it promotes the latter).

I think that once we reflect about it, this conclusion is actually what we would expect. As noted above, in the personal realm the most obvious day-to-day example of paternalistic practices is (as the name suggests) found in a parent's relationship to their child. Moreover, some of the paternalism in play in this regard is clearly epistemic in nature, as it involves the parent exercising control over the epistemic environment of the child by, for example, limiting exposure to certain sources of information. For example, if a parent thought that their child might be taken in by the joke history book that they have on their shelves, then they might well be inclined to move it out of reach.

Interestingly, the epistemic paternalist strategies employed by parents are not exclusively concerned with ensuring that their offspring gain true beliefs rather than false ones. Indeed, I take it that one very natural role for a parent to take is to try to put their child in a position where they can come to find things out for themselves, and thereby attain an understanding of the relevant subject matter. While there might be occasions where simply 'feeding' the child the truth may be the best strategy, clearly where circumstances permit it is more beneficial to the child to be fed the information she needs in order to determine the answer for themselves (e.g., by guiding them to a reliable information source and showing them how to make best use of it). The child's epistemic good is thus paternalistically furthered by its parents by their promotion not only of true belief but also of the child's capacity for understanding.²⁵

²⁵ The way in which parents monitor their child's epistemic environment has been extensively discussed in the recent literature on the epistemology of testimony. For an excellent recent discussion of this phenomenon, see Goldberg (2008). For further more general discussion of epistemic monitoring, including where it includes cases where the epistemic environment be-

What goods are being promoted by these paternalistic strategies? Clearly, there is a practical utility in children gaining true as opposed false beliefs about a certain subject matter (they will do better on exams, for example), and also to children gaining an understanding of that subject matter (where, I take it, the practical value of the latter is *ceteris paribus* typically greater than the practical value of the former). Arguably, as we have seen, there is also a specifically epistemic value involved in true belief, and which is thus also present in understanding to the extent that it involves true belief. That is, parents may well promote such goods because they (rightly) value the final goodness of truth.

But there is also a third category of good that the parents are likely promoting in pursuing these paternalistic strategies, which is the promotion of ethical goods. This is most obviously so where the epistemic paternalism concerns the promotion of understanding, since as we saw above, this is plausibly an epistemic standing which has a final value which is specifically ethical rather than epistemic, in that its goodness is a constitutive component of a life of flourishing. Indeed, I take it that most parents employ paternalistic epistemic strategies for at least partly this reason. For while parents no doubt care about the truth, and wish to instil this in their offspring, and while parents want children to have the kind of useful true beliefs that will enable them to prosper from a practical point of view, they also recognise that a good life is one in which children come to acquire a genuine understanding of a world around them (i.e., as opposed to merely having lots of true beliefs about this world).²⁶

We are now in a position to offer a more nuanced account of what epistemic paternalism might involve. First, let's sketch a scenario which might call for some variety of epistemic paternalism.

As is now well known, people are subject to various fairly systematic cognitive biases governing our judgements across a wide range of subject matters. It is also now a familiar point of social psychology that there are 'situational' elements that one can introduce to a subject's judgement context which will either neutralise or exacerbate these biases. For example, small, apparently irrelevant, changes in how a problem is posed can have a dramatic effect on whether the subject is able to resolve that problem.²⁷

ing monitored is an adult's epistemic environment, see Goldberg (2005; 2010; 2011). See also Kallestrup & Pritchard (2012; 2013).

²⁶ One can find analogous claims about the goals of education in the epistemology of education literature. See, for example, Siegel (2003), Robertson (2009), and Pritchard (*forthcoming*).

²⁷ For an excellent overview of the relevant empirical literature, see Doris (2002). See also Prinz (2009). For a comprehensive and up-to-date bibliographical survey of the literature on

Now imagine that one of these domains is practically important perhaps, say, it has a bearing on agents' financial arrangements in old age, and thus on agents' welfare, broadly conceived. If one could subtly influence the judgements made by agents in this domain without their explicit consent to ensure that they avoided cognitive biases and exercised correct judgement, then isn't there at least a *prima facie* case to do so?²⁸

With paternalism as an option on the table, we can now conceive of three distinct ways in which matters might proceed which are delineated in terms of the goals to which the paternalistic strategy is directed towards. The first is a form of paternalism which is not specifically epistemic in any interesting sense. This is where one paternalistically promotes the relevant epistemic good in this case, the promotion of correct judgements within the domain purely because this serves a non-epistemic aim, such as promoting the welfare of the subjects concerned. Here the epistemic goods promoted are merely means to a further non-epistemic end, which is why this form of paternalism isn't specifically epistemic in a substantive sense. To reflect this point, call this *quasi-epistemic paternalism*.

At the other extreme, a second kind of paternalism which clearly is specifically epistemic would involve promoting these epistemic goods purely because of their epistemic goodness. On this view, even if the agents' welfare (or indeed any further non-epistemic good) weren't furthered by this paternalistic strategy, so long as the epistemic goods are enhanced to a degree compatible with other values e.g., so long as this paternalistic strategy did not unduly undermine the subjects' autonomy then the strategy would be pursued. Call this *robust epistemic paternalism*.

A less extreme form of epistemic paternalism, though still plausibly a distinctively epistemic form of paternalism, would be where one paternalistically promotes the epistemic goods purely because of their goodness (i.e., and not merely as a means to promote some other goods), but

situationism and its relevance to virtue theory (including both empirical and philosophical literature, and covering both historical and contemporary sources), see Alfano & Fairweather (*forthcoming*). For discussion of the some of the epistemological implications of situationism, see Axtell (2010), Baehr (2011, *passim*), Alfano (2012; 2013, part II; *forthcoming*), Hazlett (2013, *passim*), Doris & Olin (*forthcoming*), Fairweather & Montemayor (*forthcoming*), and Pritchard (*forthcoming*).

²⁸ This is roughly what is proposed by a so-called 'nudge' approach to public policy, as expressed in Thaler & Sunstein (2008). Interestingly, this methodology has been endorsed by a number of governments, including the current UK government, which has its own dedicated 'nudge' public policy unit (called the 'Behavioural Insights Team'). For discussion of this approach to public policy, see Haynes *et al* (2012).

where this goodness is not restricted to a specifically epistemic kind of goodness. So, for example, one might think it desirable to paternalistically promote the epistemic good within this domain even if it did not generate any wider benefits, such as in terms of the agents' welfare. Crucially, however, and unlike robust epistemic paternalism, the rationale for promoting the epistemic good would not be because of its exclusively epistemic value, but rather because the epistemic good was also of value in non-epistemic ways. Perhaps, for example, one holds that by promoting correct judgement in this domain the subjects concerned will attain epistemic goods like understanding which are valuable in their own right in virtue of their non-epistemic goodness (e.g., in virtue of an essential role they play in a life of flourishing). Call this *modest epistemic paternalism*.

The distinction between modest epistemic paternalism and *quasi*-epistemic paternalism is marked by whether the paternalism in question is ultimately motivated by the goal of promoting an epistemic good. Whereas *quasi*-epistemic paternalism is only concerned with the goal of promoting an epistemic good as a means of promoting a further non-epistemic good, modest epistemic paternalism is focussed on promoting the epistemic good in its own right. Unlike robust epistemic paternalism, however, modest epistemic paternalism is not solely concerned with promoting a specifically epistemic kind of goodness. The distinction between robust and modest epistemic paternalism thus reflects the distinction drawn earlier regarding epistemic value and the value of the epistemic.

We can set *quasi*-epistemic paternalism to one side for our purposes, since ultimately it does not raise questions about the merits of paternalistic strategies which are specific to epistemology. Let us then focus on modest and robust epistemic paternalism.

The latter will be the most controversial. Its plausibility will depend on whether epistemic goodness is good *simpliciter* i.e., whether it is the kind of thing that can plausibly go into the mix when determining all things considered goodness. For if the epistemic good is not good *simpliciter*, then it is hard to see why we would promote it at the expense of individual freedom and autonomy (which clearly are valuable *simpliciter*). We noted earlier that a *prima facie* case could be made for the goodness *simpliciter* of the epistemic, and to the extent that this case succeeds there is therefore a rationale for some degree of robust epistemic paternalism. Even then, however, given the relative weight accorded to such values as individual freedom and autonomy, there will be a challenge in play to demonstrate that the benefits on offer in terms of the epistemic good suffice to outweigh the disbenefits that such a strategy will produce.

The case for modest epistemic paternalism is much easier to make. We've seen above that plausibly there are epistemic standings, such as understanding, which have a value which extends beyond the epistemic. Moreover, this non-epistemic value that they possess is an even more credible candidate to be a kind of goodness which is good *simpliciter*. So long as the paternalistic strategy in play does not merely promote mere true belief, but also further epistemic standings, such as understanding, then it can in principle be defended in terms of how it generates epistemic goods which are both epistemically valuable and non-epistemically valuable, and where the goodness generated is at least in part goodness *simpliciter*. Even here, the challenge will be to demonstrate that the potential benefits outweigh the potential costs, but clearly the task in hand will be much easier than it was for robust epistemic paternalism.

5. *Concluding remarks*

I have argued that a more nuanced conception of epistemic goodness leads to a more subtle conception of epistemic paternalism. In so doing, I have shown that evaluating the prospects for epistemic paternalism is a more complex matter than we might have hitherto supposed. Interestingly, however, even despite this complexity we have also seen that there are genuine forms of epistemic paternalism which are in principle defensible. I want to close by noting two further complications and issues which are raised by the conception of epistemic value and epistemic paternalism offered here.

Consider first the problems that might be raised by competing epistemic goods. For example, a strategy that maximises an agent's true beliefs might actually *undermine* her prospects for gaining understanding. This might occur in a case where one obscures some relevant feature of the agent's environment because one is aware that if she were to attend to it then this would cause her to doubt some true propositions which she would otherwise believe. Even so, attending to this feature of her environment could be a crucial step towards gaining a proper understanding of her cognitive situation. By the same token, a strategy aimed at promoting the subject's understanding in such a case could actually undermine her prospects for gaining true beliefs.

Given that the potential for conflicting epistemic goods in this sense, this will clearly have implications for epistemic paternalist strategies. In particular, robust and modest epistemic paternalist strategies may well di-

verge quite considerably in terms of what they aim to achieve. So, for example, if one is pursuing a robust epistemic paternalism, and one follows veritism in taking the ultimate epistemic good to be true belief, then in the case just described one would presumably be content to promote the epistemic good of true belief over that of understanding.²⁹ In contrast, if one is pursuing a modest epistemic paternalism, and one holds that understanding is an epistemic good with a distinctive non-epistemic value, then one would presumably opt for a strategy which focuses on the promotion of understanding even at the expense of maximising true belief. If one holds that epistemic standings are valuable in a range of different ways, then the picture that emerges could potentially be even more complex, as one has to trade-off the promotion of some epistemic goods over others.

A second issue that should be borne in mind when we consider epistemic paternalism is the extent to which paternalism of this stripe could in the longer term be conducive to the subject's freedom and autonomy. In normal cases of paternalism one limits the subject's freedom and autonomy in order to promote wider ends, such as their welfare. As we have seen, the same limitation of individual freedom and autonomy would occur with epistemic paternalism. Interestingly, however, a form of epistemic paternalism which promoted epistemic goods like knowledge and understanding might well serve to *promote* the subject's freedom and autonomy over the longer term. After all, having a correct understanding of one's situation and of one's options can have a substantial bearing on one's ability to exercise one's freedom, and to take cognitive ownership of one's decisions in the manner of an autonomous individual. This means that epistemic paternalist strategies which promote epistemic goods could potentially thereby contribute to an agent's freedom and autonomy over the longer term, even while undermining freedom and autonomy over the short term. Where this is the case, it clearly will have a bearing on how we evaluate the legitimacy of these strategies.³⁰

²⁹ I say 'presumably' because one option open to the veritist albeit an option not normally recognised as being open so far as much of mainstream epistemology is concerned is to maintain that the goal of maximizing true belief should not be understood in terms of simply aiming for the largest number of true beliefs on the part of the subject. For two excellent discussions of this point, see Treanor (2012, *forthcoming*). See also Pritchard (*forthcoming*).

³⁰ This paper was written while I was in receipt of a Phillip Leverhulme Prize, and I am grateful to them for their support. An earlier version of this paper was presented at the *Epistemic Inclusiveness and Trust* conference hosted by the *Social Epistemology Research Group* at the University of Copenhagen in August 2011, and I am grateful to the audience on this occasion for their feedback. Thanks also to Kristoffer Ahlstrom, Alvin Goldman, Klemens Kappel, Hilary Kornblith, Erik Olsson, Christian Piller, Mike Ridge, Shane Ryan and Ernest Sosa.

References

- Ahlstrom, Kristoffer H., 2011, "Epistemic Paternalism: A Defence", *manuscript*.
- Alfano, Mark, 2012, "Expanding the Situationist Challenge to Responsibility Virtue Epistemology", in *Philosophical Quarterly*, 62: 223-249.
- Alfano, Mark, 2013, *Character as Moral Fiction*, Cambridge University Press, Cambridge.
- Alfano, Mark, *Forthcoming*, "Extending the Situationist Challenge to Reliabilism about Inference", in *Virtue Scientia: Bridges Between Philosophy of Science and Virtue Epistemology*, ed. A. Fairweather, Springer, Dordrecht, Holland.
- Alfano, Mark, Fairweather, Abrol, *Forthcoming*, "Situationism and Virtue Theory", in *Oxford Bibliographies: Philosophy*.
- Axtell, Guy, 2010, "Agency Ascriptions in Ethics and Epistemology: Or, Navigating Intersections, Narrow and Broad", in *Metaphilosophy*, 41: 73-94.
- Baehr, Jason, 2011, *The Inquiring Mind: On Intellectual Virtues and Virtue Epistemology*, Oxford University Press, Oxford.
- Brady, Michael S., Pritchard, Duncan H., 2003, *Moral and Epistemic Virtues*, eds., Blackwell, Oxford.
- David, Marian, 2001, "Truth as the Epistemic Goal", in *Knowledge, Truth, and Duty: Essays on Epistemic Justification, Virtue, and Responsibility*, ed. M. Steup, 151-169, Oxford University Press, Oxford.
- David, Marian, 2005, "Truth as the Primary Epistemic Goal: A Working Hypothesis", in *Contemporary Debates in Epistemology*, eds. E. Sosa & M. Steup, 296-312, Blackwell, Oxford.
- Doris, John, 2002, *Lack of Character: Personality and Moral Behaviour*, Cambridge University Press, Cambridge.
- Doris, John, Olin, Lauren, *Forthcoming*, "Vicious Minds: Virtue, Cognition, and Skepticism", in *Philosophical Studies*.
- Dworkin, Gerald, 2010, "Paternalism", in *Stanford Encyclopaedia of Philosophy*, ed. E. Zalta, <http://plato.stanford.edu/entries/paternalism/>.
- Elgin, Catherine, 1996, *Considered Judgement*, Princeton University Press, Princeton, NJ.
- Elgin, Catherine, 2004, "True Enough", in *Philosophical Issues*, 14: 113-131.
- Elgin, Catherine, 2009, "Is Understanding Factive?", in *Epistemic Value*, eds. A. Haddock, A. Millar & D. H. Pritchard, 322-330, Oxford University Press, Oxford.
- Fairweather, Abrol, Montemayor, Carlos, *Forthcoming*, "The Meeting of the Twain: Alfano on Situationism and Virtue Epistemology", in *Philosophical Quarterly*.

- Feinberg, Joel, 1986, *Harm to Self*, Oxford University Press, Oxford.
- Geach, Peter T., 1956, "Good and Evil", in *Analysis*, 17: 32-42.
- Goldberg, Sanford, 2005, "Testimonial Knowledge from Unsafe Testimony", in *Analysis*, 65: 302-311.
- Goldberg, Sanford, 2008, "Testimonial Knowledge in Early Childhood, Revisited", in *Philosophy and Phenomenological Research*, 76, 1-36.
- Goldberg, Sanford, 2010, *Relying on Others: An Essay in Epistemology*, Oxford University Press, Oxford.
- Goldberg, Sanford, 2011, "The Division of Epistemic Labour", in *Episteme*, 8: 112-125.
- Goldman, Alvin, 1991, "Epistemic Paternalism: Communication Control in Law and Society", in *Journal of Philosophy*, 88, 113-131.
- Goldman, Alvin, 1999, *Knowledge in a Social World*, Oxford University Press, Oxford.
- Goldman, Alvin, 2002, "The Unity of the Epistemic Virtues", in *Pathways to Knowledge: Private and Public*, Oxford University Press, Oxford: 51-72.
- Goldman, Alvin, Olsson, Erik J., 2009, "Reliabilism and the Value of Knowledge", in *Epistemic Value*, eds. A. Haddock, A. Millar & D. H. Pritchard, Oxford University Press, Oxford: 19-41.
- Greco, John, 2003, "Knowledge as Credit for True Belief", in *Intellectual Virtue: Perspectives from Ethics and Epistemology*, eds. M. DePaul & L. Zagzebski, Oxford University Press, Oxford: 111-134.
- Greco, John, 2007, "The Nature of Ability and the Purpose of Knowledge", in *Philosophical Issues*, 17: 57-69.
- Greco, John, 2008, "What's Wrong With Contextualism?", in *Philosophical Quarterly*, 58: 416-436.
- Greco, John, 2009a, *Achieving Knowledge*, Cambridge University Press, Cambridge.
- Greco, John, 2009b, "The Value Problem", in *Epistemic Value*, eds. A. Haddock, A. Millar & D. H. Pritchard, Oxford University Press, Oxford: 313-321.
- Greco, John, 2009c, "Knowledge and Success From Ability", in *Philosophical Studies*, 142: 17-26.
- Hazlett, Allan, 2013, *A Luxury of the Understanding: On the Value of True Belief*, Oxford University Press, Oxford.
- Haynes, Laura, Service, Owain, Goldacre, Ben, Torgerson, David, 2012, "Test, Learn, Adapt: Developing Public Policy with Randomised Controlled Trials", UK Cabinet Office Policy Paper, www.gov.uk/government/publications/test-learn-adapt-developing-public-policy-with-randomised-controlled-trials.

- Jones, Ward, 1997, "Why Do We Value Knowledge?", in *American Philosophical Quarterly*, 34: 423-440.
- Kallestrup, Jesper, Pritchard, Duncan H., 2012, "Robust Virtue Epistemology and Epistemic Anti-Individualism", in *Pacific Philosophical Quarterly*, 93: 84-103.
- Kallestrup, Jesper, Pritchard, Duncan H., 2013, "Robust Virtue Epistemology and Epistemic Dependence", in *Knowledge, Virtue and Action: Putting Epistemic Virtues to Work*, eds. T. Henning & D. Schweikard, Routledge, London, ch. 11.
- Kleinig, John, 1983, *Paternalism*, Rowman and Allanheld, Totowa, NJ.
- Kvanvig, Jonathan, 2003, *The Value of Knowledge and the Pursuit of Understanding*, Cambridge University Press, Cambridge.
- Kvanvig, Jonathan, 2005, "Truth is not the Primary Epistemic Goal", in *Contemporary Debates in Epistemology*, eds. E. Sosa & M. Steup, Blackwell, Oxford: 285-296.
- Olsson, Erik J., 2007, "Reliabilism, Stability, and the Value of Knowledge", in *American Philosophical Quarterly*, 44: 343-355.
- Olsson, Erik J., 2009, "In Defence of the Conditional Probability Solution to the Swamping Problem", in *Grazer Philosophische Studien*, 79: 93-114.
- Percival, Philip, 2003, "The Pursuit of Epistemic Good", in *Metaphilosophy*, 34: 29-47; and reprinted in Brady & Pritchard, 2003: 29-46.
- Prinz, Jesse, 2009, "The Normativity Challenge: Cultural Psychology Provides the Real Threat to Virtue Ethics", in *Journal of Ethics*, 13: 117-144.
- Pritchard, Duncan H., 2007, "Recent Work on Epistemic Value", in *American Philosophical Quarterly*, 44: 85-110.
- Pritchard, Duncan H., 2009a, "Knowledge, Understanding and Epistemic Value", in *Epistemology (Royal Institute of Philosophy Lectures)*, ed. A. O'Hear, Cambridge University Press, Cambridge: 19-43.
- Pritchard, Duncan H., 2009b, "The Value of Knowledge", in *Harvard Review of Philosophy*, 16: 2-19.
- Pritchard, Duncan H., 2010, "Achievements, Luck and Value", in *Think*, 25: 1-12.
- Pritchard, Duncan H., 2011, "What is the Swamping Problem?", in *Reasons for Belief*, eds. A. Reisner & A. Steglich-Petersen, Cambridge University Press, Cambridge: 244-259.
- Pritchard, Duncan H., 2012, "Anti-Luck Virtue Epistemology", in *Journal of Philosophy*, 109: 247-279.
- Pritchard, Duncan H., *Forthcominga*, "Epistemic Virtue and the Epistemology of Education", in *Journal of Philosophy of Education*; to be reprinted in *Education and the Growth of Knowledge: Perspectives From Social and Virtue Epistemology*, ed. B. Kotzee, Blackwell, Oxford.

- Pritchard, Duncan H., *Forthcomingb*, “Re-evaluating the Epistemic Situationist Challenge to Virtue Epistemology”, in *Naturalizing Epistemic Virtue*, eds. A. Fairweather & O. Flanagan, Cambridge University Press, Cambridge.
- Pritchard, Duncan H., *Forthcomingc*, “Truth as the Fundamental Epistemic Good”, in *The Ethics of Belief: Individual and Social*, eds. J. Matheson & R. Vitz, Oxford University Press, Oxford.
- Pritchard, Duncan H., *Forthcomingd*, “Veritism and Epistemic Value”, in *Alvin Goldman and His Critics*, eds. H. Kornblith & B. McLaughlin, Blackwell, Oxford.
- Pritchard, Duncan H., Millar, Alain, Haddock, Adrian, 2010, *The Nature and Value of Knowledge: Three Investigations*, Oxford University Press, Oxford.
- Pritchard, Duncan H., Turri, John, 2011, “Knowledge, the Value of”, in *Stanford Encyclopaedia of Philosophy*, ed. E. Zalta, <http://plato.stanford.edu/entries/knowledge-value/>.
- Riggs, Wayne, 2002a, “Beyond Truth and Falsehood: The *Real* Value of Knowing that *P*”, in *Philosophical Studies*, 107: 87-108.
- Riggs, Wayne, 2002b, “Reliability and the Value of Knowledge”, in *Philosophy and Phenomenological Research*, 64: 79-96.
- Riggs, Wayne, 2008, “The Value Turn in Epistemology”, in *New Waves in Epistemology*, eds. V.F. Hendricks & D.H. Pritchard, Palgrave Macmillan, Basingstoke, UK: 300-323.
- Riggs, Wayne, 2009, “Getting the *Meno* Requirement Right”, in *Epistemic Value*, eds. A. Haddock, A. Millar & D. H. Pritchard, Oxford University Press, Oxford: 331-338.
- Robertson, Emily, 2009, “The Epistemic Aims of Education”, in *Oxford Handbook of Philosophy of Education*, ed. H. Siegel, Oxford University Press, Oxford: 11-34.
- Sartorius, Rolf, ed., 1983, *Paternalism*, University of Minnesota Press, Minneapolis, MN.
- Siegel, Harvey, 2003, “Cultivating Reason”, in *Companion to the Philosophy of Education*, ed. R. Curren, Blackwell, Oxford: 305-319.
- Sosa, Ernest, 1988, “Beyond Skepticism, to the Best of Our Knowledge”, in *Mind*, 97: 153-189.
- Sosa, Ernest, 1991, *Knowledge in Perspective: Selected Essays in Epistemology*, Cambridge University Press, Cambridge.
- Sosa, Ernest, 2007, *A Virtue Epistemology: Apt Belief and Reflective Knowledge*, Clarendon Press, Oxford.
- Sosa, Ernest, 2009a, *Reflective Knowledge: Apt Belief and Reflective Knowledge*, Clarendon Press, Oxford.

- Sosa, Ernest, 2009*b*, "Situations Against Virtues: The Situationist Attack on Virtue Theory", in *Philosophy of the Social Sciences: Philosophical Theory and Scientific Practice*, ed. C. Mantzavinos, Cambridge University Press, Cambridge: 274-290.
- Swinburne, Richard, 1999, *Providence and the Problem of Evil*, Oxford University Press, Oxford.
- Thaler, Richard, Sunstein, Cass, 2008, *Nudge: Improving Decisions About Health, Wealth, and Happiness*, Yale University Press, Ithaca, NJ.
- Treanor, Nick, 2012, "The Measure of Knowledge", in *Noûs*.
- Treanor, Nick, *Forthcominga*, "Trivial Truths and the Aim of Inquiry", in *Philosophy and Phenomenological Research*.
- Van DeVeer, Donald, 1986, *Paternalistic Intervention*, Princeton University Press, Princeton, NJ.
- Williamson, Timothy, 2000, *Knowledge and Its Limits*, Oxford University Press, Oxford.
- Zagzebski, Linda, 1996, *Virtues of the Mind: An Inquiry into the Nature of Virtue and the Ethical Foundations of Knowledge*, Cambridge University Press, Cambridge.
- Zagzebski, Linda, 1999, "What is Knowledge?", in *The Blackwell Guide to Epistemology*, eds. J. Greco & E. Sosa, Blackwell, Oxford: 92-116.
- Zagzebski, Linda, 2001, "Recovering Understanding", in *Knowledge, Truth and Duty: Essays on Epistemic Justification, Responsibility and Virtue*, ed. M. Steup, Oxford University Press, Oxford: 235-252.
- Zagzebski, Linda, 2003, "The Search for the Source of the Epistemic Good", in *Metaphilosophy*, 34: 12-28; and reprinted in Brady & Pritchard, 2003: 13-28.