

MUST WE WORRY ABOUT EPISTEMIC SHIRKERS?

Daniele Bruno (Humboldt-Universität zu Berlin)

daniele.bruno@hu-berlin.de

This is the penultimate submitted manuscript version of the paper published in *Inquiry*.

To cite, please refer to the published version:

<https://doi.org/10.1080/0020174X.2024.2338791>

Abstract

It is commonly assumed that blameworthiness is epistemically constrained. If one lacks sufficient epistemic access to the fact that some action harms another, then one cannot be blamed for harming. Acceptance of an epistemic condition for blameworthiness can give rise to a worry, however: could agents ever successfully evade blameworthiness by deliberately stunting their epistemic position? I discuss a particularly worrisome version of such epistemic shirking, in which agents preemptively seek to avoid access to potentially morally relevant facts. As Roy Sorensen and Jan Willem Wieland have argued, we seem to be faced with a potentially troubling regress when trying to explain what goes wrong in such situations. I argue that the solution to this so-called Shirker Problem is not to be found in complicated and demanding anti-shirking-principles with universal scope and potentially self-reflexive content. Instead, careful consideration of the necessary motivational make-up of genuine shirkers reveals their ignorance to lack the necessary depth to be exculpatory. Since shirkers by definition react to the possibility of the revelation of morally relevant facts with acts of epistemic self-limitation, they necessarily prove guilty of a form of culpable moral recklessness. Epistemic shirking thus does not pose a worrisome problem for ethical theory.

1) Introduction

Sometimes, people do things we would normally consider to be morally wrong, yet we correctly excuse them because they were (to some extent) ignorant of the import of their actions. If I never had the chance to find out about my best friend's recent spell of depression (she is very good at hiding her feelings), it would not be appropriate to blame me for not assisting her in getting help. Accepting any sort of epistemic condition on blameworthiness, however, might give rise to a worry: if blameworthiness is irretrievably linked with conditions of knowability, what about agents who consciously *embrace* or *seek* ignorance on matters of moral relevance? Given the control that we have over many aspects of our own epistemic situations, it seems quite feasible for us to make access to some morally relevant fact impossible for ourselves. As I will show in what follows, this

phenomenon is in fact widespread, more than it might seem at first glance. But does this possibility then afford agents the opportunity to actively create excuses for themselves, effectively evading blameworthiness for failing to meet some of their moral obligations?

In this paper I will have a closer look at phenomena of this kind, which I shall collectively call cases of *wilful ignorance*.¹ This tag unites a host of issues, corresponding to a wide range of behaviours and psychological phenomena. The philosophically most interesting cases, I will argue, are ones involving what I shall call *epistemic shirking* (following Wieland 2015). Epistemic shirkers are agents who deliberately seek to *pre-empt* the possibility of their coming to know that they have an obligation, with the aim of eliminating their blameworthiness for what would otherwise have constituted a blameworthy violation of this obligation. These cases are of particular interest as they seem to give rise to a troubling infinite regress, a problem that was first discussed in a debate between Ted Sider (1995) and Roy Sorensen (1995), and more recently again brought to the forefront by Jan Willem Wieland (2015).

After laying out this ‘Shirker Problem’ in sections 3 and 4, I will present Wieland’s ‘strong’ solution involving self-referential obligations in section 5 and explain why I find it unsatisfactory. In section 6, I will put forward a theoretically less ambitious solution, aiming to show that the very idea of shirking on an obligation implies a blameworthy disrespect of the respective obligation on the part of the shirker – specifically given her risk of breaching it. Therefore, or so I shall argue, in every instance of epistemic shirking we can pinpoint a risk-based moral obligation that the shirker impermissibly disregards.

2) The Epistemic Condition for Blameworthiness

As mentioned above, the starting point of the discussion on wilful ignorance is the intuitively highly plausible idea that ignorance of morally relevant facts can sometimes render blame unjustified. This idea is reflected in the following principle:

Epistemic condition (EC): For all putative obligations x , S is blameworthy for violating x only if S is in a position to know (with reasonable effort) that they are subject to x .

Note that EC, as formulated, leaves open the question of what exactly the grounds of ascribing blamelessness to a sufficiently ignorant agent are. There are two principal options available here. Either S fails to be blameworthy in the relevant situations because S is not appropriately held morally responsible for her nonetheless morally wrong

¹ Wilful ignorance (cf. also Lynch 2016, Wieland 2017b, Sarch 2018) also passes under some different labels: Wilful blindness (Simon 2005), affected ignorance (Moody-Adams 1994, N. Williams 2008) or strategic ignorance (Wieland 2017a), for example.

behaviour, or S cannot even be subject to x if S is not in a position to know so. On this latter view, our moral obligations themselves are subject to epistemic restrictions, either because what we are obligated to do depends only on things that we necessarily have epistemic access to (our mental states, for example), or because there is what we may call an *epistemic filter* on obligations, such that only facts epistemically accessible to a given agent can impinge on what this agent (morally) ought to do.²

As stated, EC is of course only a very rough formulation of an epistemic condition for blameworthiness. Several of its elements would need to be specified further in order to yield anything close to concrete answers to the question of which epistemic situations can support ascriptions of blame to agents. Most importantly, one would need to precisify what it means to be in a position to know a certain proposition, and what kind of epistemic access is thus necessary for blameworthiness. There are two dimensions that are of principal importance here. For one thing, one would have to settle on the demandingness of the ability ascription in question. One can imagine a variety of levels of difficulty of acquiring knowledge that p, which can range from easy (making a simple inference, recalling past experiences, reading a text laid out before one's eyes) through medium (doing basic research on the internet, connecting various bits of knowledge to draw a non-straightforward inference) to very hard (obtaining a lost ancient document, solving yet unsolved complex mathematical equations). While the hardest cases are clearly excluded, where exactly to draw the line of the relevant notion of knowability will be a matter of controversy. On the other hand, there is also an important variance regarding the costs of acquiring knowledge for the agent. These perhaps most crucially include the time invested in inquiry, but may also be of a different sort.³ Imagine that someone credibly relates to you that they have a document containing important information relevant to your moral obligations (perhaps it is the fact that your best friend is depressed and in need of help). But they will only part with this document in exchange for all of your belongings. Although it may not be strictly speaking *difficult* for you to obtain the information in the document in this case, you still seem blameless for failing to obtain it, and therefore for failing to help your friend. The reason for this is quite simply that we cannot reasonably expect you to incur these costs for the sake of coming to know your obligations. Because of this, EC must contain a reasonability clause ('with reasonable effort') in order to be plausible. What exact degree of both cost and difficulty we can expect an agent to bear to obtain morally relevant information before it becomes inappropriate to blame them is perhaps *the* most important question pertaining to EC that I will have to leave open in the context of this paper.

² Versions of this kind of *Perspectivism about Ought* are defended by Lord (2015), Kieseewetter (2017), and Andrić (2021)

³ Cf. also D. Williams 2021: 7809.

Before I move on, let me add one further important caveat. As I already have mentioned, my discussion will proceed under the (widely shared) assumption that blameless ignorance often does excuse. There is, however, a heated debate about which *kind* of ignorance can excuse. While some philosophers hold that actions can legitimately be excused both by ignorance of descriptive truths (e.g. that your friend is depressed) and by ignorance of moral truths (e.g. that one morally ought to help one's friends when one is the only person in a position to do so), others hold that moral ignorance never excuses.⁴ Since endorsing the latter view would provide a short-cut around many (though not all) of the more complicated problems about wilful ignorance, I shall assume for the sake of the discussion that being ignorant of one's moral obligations can at least sometimes excuse an agent.

3) Wilful Ignorance and Epistemic Shirking

Let us now return to wilful ignorance. Though I will not offer a definition of wilful ignorance, it will nonetheless be useful for my discussion to at least give a brief overview over the wide range of phenomena that philosophers have discussed under this heading. This will also allow me to more clearly delineate the specific phenomenon of epistemic shirking that will be the focus of my discussion.

One contribution that was integral to bringing the importance of wilful ignorance to the forefront of philosophical attention is Michele Moody-Adams' excellent 'Culture, Responsibility, and Affected Ignorance'. In it, she deals with the phenomenon of what she calls *affected ignorance*, or, as she puts it 'choosing not to know what one can and should know' (Moody-Adams 1994: 296), a phenomenon that Moody-Adams claims to be at the heart of many of history's most pervasive wrongs and injustices.

Affected ignorance, as understood by Moody-Adams, is a highly gerrymandered phenomenon, which manifests itself in several different ways. To name just a few examples brought up by her, it might involve people refusing to acknowledge the connection between their actions and the consequent suffering of their victims (affecting a psychological disconnect between one's knowledge of the methods of industrial animal agriculture and the consumption of its products), linguistic sugar-coating ('enhanced interrogation' instead of 'torture') or simple instances of 'looking the other way'. As Moody-Adams and others convincingly argue, all these types of affected ignorance constitute a serious *practical* problem – they perpetuate cultures of injustice and wrongdoing by rendering public debate of the problematic practices more difficult, as well as by obstructing our ability to accurately judge individuals in matters of responsib-

⁴ See e.g. Harman 2011, 2015.

ility for their culturally sanctioned wrongdoing.⁵ It has been argued that philosophical engagement with affected ignorance is thus of great importance if one wants to understand and tackle endemic and systematic injustices such as racism (Mills 2017, Martín 2021), sexism (Calhoun 1989, Medina 2013, Code 2020), or the abuse of animals in factory farming (N. Williams 2008).

However, few of these situations seem to pose a serious *theoretical* problem from the point of view of the ethical theorist concerned with the implications of the epistemic condition. In many of the situations described, the ignorance of wrongdoing is quite literally *affected*, that is, at least partly not natural or genuine, but feigned. The agents in question therefore do not actually fail to meet the conditions of EC.⁶ Of course, most agents engaging in affected ignorance may not really seek actual avoidance of the conditions of EC – they may wish to simply make sure they do not become targets of others’ blame or sanctions, or to evade the gnawing feeling of guilt that comes with awareness of one’s own moral failings. As such, it may be enough for them to simply engage in a thorough kind of self- and other-deception. Of course, this kind of deception might ultimately become so firmly engrained, both individually and socially, as to become entirely unnoticeable even from the first-person perspective. Since there is this important phenomenon of initially feigned or at least less-than-fully-real ignorance, I prefer deviating from Moody-Adam’s own usage and retaining the use of the term ‘affected ignorance’ for only this phenomenon. I will instead use the wider term ‘wilful ignorance’ for the more general phenomenon of ‘choosing not to know what one can and should know’.

Understood in this wider sense, wilful ignorance does not necessarily involve any self- or other-deception.⁷ Many types of actions that people can and regularly do take lead to an *actual* deterioration of their epistemic circumstances, yielding a situation in which they in fact are no longer in a position to know certain propositions. It is with this specific kind of wilful ignorance that I will be concerned in what follows. My focus will be on agents who not only wish to evade blame or feelings of guilt by assuming a feigned or irregular ignorance, but rather those who attempt to make a difference to their *blameworthiness* by ensuring that they do not get to form the respective doxastic attitudes in the first place, usually by affecting their own evidential situations. In the following, I will refer to any attempts to change the moral situation by causing self-effected ignorance as *shirking*⁸:

⁵ For the political importance of wilful ignorance, see also D. Williams 2021.

⁶ This is of course precisely what Moody-Adams argues for. Her principal aim is challenging widespread assumptions about the excusing force of cultural membership in societies in which certain unjust practices are so firmly ingrained that ignorance of their injustice borders on universal. In many, if not all of these societies, Moody-Adams argues, most of the ignorance is affected in one way or another, which in turn means that the members are not excused on epistemic grounds.

⁷ For an extensive argument defending this thesis, see also Lynch 2015.

⁸ In this nomenclature I follow Jan Willem Wieland (2015).

Shirking: An agent S shirks on an obligation to ϕ iff S attempts to evade blameworthiness for violating an obligation to ϕ at t_2 by taking an action A at t_1 , which will cause S to not meet the necessary conditions for being blameworthy for (wrongly) not ϕ -ing at t_2 .

There are two principal ways in which an agent can attempt to evade blameworthiness for failing to perform an action: either by affecting their control or their good epistemic standing. Possession of both is, so we are assuming, necessary for blameworthiness.

In the *control case*, the agent purposefully puts themselves in a situation in which they are no longer able to discharge the obligation they know to have. Imagine someone, call him Leroy, who has promised to be present at a party on Friday, but does not want to go. If on Thursday, Leroy locks himself in his basement and flushes away the key, he makes it the case that on Friday, it is no longer within his control to attend.

In the *epistemic case*, the agent makes sure that at t_2 , they will not (or no longer) be in a position to know about their obligation to ϕ . Instead of locking himself up, for example, Leroy could also take an amnesia pill that leads him to forget that he ever uttered the promise. Either way, at t_2 some necessary condition for blameworthiness will not be met by him.

At first glance, it thus might seem that we are committed to accept that Leroy is not blameworthy for no-showing at the party. But of course, agents attempting to evade blameworthiness will not get off the hook that easily. We can still justify ascriptions of blame at t_2 through a process commonly called *tracing*. Leroy's blameworthiness for not showing up at the party can reasonably be considered derivative of his blameworthiness for committing the prior wrong of locking himself up or taking the pill. Here, I will not go into the intricacies of the debate on tracing, which is itself complicated and controversial.⁹ For our purposes, it is enough to note the following: in order to justify blameworthiness through the process of tracing, it is necessary that one is able to point to an earlier wrong act, sometimes called the *benighting act* (H. Smith 1983: 547), for which the agent was blameworthy, and which stands in the right relation to the latter wrongdoing.¹⁰

⁹ Though for a thoroughgoing defence of tracing, see Fischer and Tognazzini 2009.

¹⁰ In the control case, we can explain the wrongness of the benighting act through the independently plausible principle of *Necessary Means Transmission for Ought* (Gertken and Kiesewetter 2021):

Necessary Means Transmission for Ought: If S ought to ϕ , and ψ -ing is a necessary means for S to ϕ , then S ought to ψ .

The action of [not throwing away the key] is a necessary means for Leroy to make it to the party: all the possible courses of action he can take that will lead him to attend include [not throwing away the key]. Therefore, if he ought to go to the party, by Necessary Means Transmission he also ought to [not throw away the key]. One interesting and tricky difference between the control case and the epistemic case is that this principle is not equally available in the latter. Knowing that he made the promise is not strictly speaking a necessary condition for Leroy showing up at the party. To explain the wrongness of epistemic shirking, other rationales are therefore required.

The Shirker Problem, first raised in a debate between Sorensen (1995) and Sider (1995) arises because of a special difficulty in pin-pointing such a benighting act in the case of epistemic shirking. For any tracing explanation to be successful, it is crucial that the benighting act itself is performed under conditions of blameworthiness. Showing that this blameworthiness obtains, however, is rendered difficult by a feature plausibly unique to epistemic shirking. Theoretically, it seems possible for agents to *pre-emptively* avoid incurring obligations by purposefully deteriorating their epistemic position in advance. Crucially, this appears to allow them to shirk obligations without learning of the nature of any obligation shirked, including obligations not to perform the act of shirking itself. This, as I will shortly show, can seem to lead to a vicious regress. First, however, let me briefly say a few words about the phenomenon of such pre-emptive epistemic shirking, both to clarify it, and to defend its very possibility.

At first sight, genuine pre-emptive epistemic shirking might seem like a relatively exotic phenomenon that we will rarely find exemplified in real-life cases.¹¹ However, I think there is a good case to be made for the claim that epistemic shirking is significantly more common than one might think. Though the renewed philosophical interest in the phenomenon is relatively recent, wilful ignorance and epistemic shirking have long been the focus of a lively debate within jurisprudence.¹² The legal theorist William H. Simon offers the following real-life example involving some of the main actors from the Watergate scandal to illustrate the phenomenon:

When Hugh Sloan, treasurer of the Committee to Re-Elect the President ("CREEP"), asked campaign finance chairman Maurice Stans why Gordon Liddy had asked for \$83,000, Stans replied "I do not want to know and you do not want to know ... (Liddy was planning the burglary of the Democratic National Committee headquarters.) When Liddy's accomplice Howard Hunt phoned presidential advisor Chuck Colson after the burglary, Colson "repeatedly insisted that he knew nothing about Watergate and wanted to keep it that way. [...] When asked much later about charges that the Greek military dictatorship made secret contributions to Watergate-related activities, CIA Director Richard Helms replied, "Even if somebody suggests they would like to do it, I would insist that they don't tell me about it because that is dynamite." (Simon 2005, p. 4)

What we see in this example is a wide array of actors who not only fail to seek further information, but actively discourage others from sharing information that they otherwise might have shared. Their aim clearly is to be in a worse evidential position than they would have been without their exhortations. What is more, this is plainly with a view to avoiding accountability. Though the actors involved in Watergate will have likely been

¹¹ For reasons of economy of language, I will drop the 'pre-emptive' modifier in what follows and simply speak about 'epistemic shirking'. The cases I am interested in, however, are exclusively of a pre-emptive nature (and thus unlike the amnesia pill case).

¹² For an overview, see Sarch 2018.

primarily concerned with avoiding *legal* or *political* accountability, we can easily imagine a case in which their concern instead lies with moral blameworthiness instead.

In the following, I will focus my attention on a fictionalised case, which will allow us to draw out the crucial elements of epistemic shirking in the clearest possible manner:

UNSUBSCRIBE: Terry is a big fan of a famous pop singer. He follows her on various social medial channels in order to stay up to date about her personal and professional developments. Recently, this singer has taken a great interest in morally motivated activism and used her platform to advocate it. Terry, however, is comfortable with his current lifestyle and would rather not find out how many of his routine behaviours are morally problematic. Though he regrets missing out on information about his favourite singer, he unsubscribes from her channels on all his social media accounts. He hopes that by not finding out more of the information the singer shares, he will not only spare himself a bad conscience, but also protect himself from blameworthiness for his future actions. As it turns out, the next day, the singer shares information about the use of child labour in the manufacturing of a certain brand of trainers. One week later, Terry buys a pair of these trainers, not knowing that child labour went into their production. Had he not decided to unsubscribe, he would have known about this fact, and furthermore known that the fact presents a strong moral reason against buying these shoes.¹³

Let's assume that at t_2 , Terry no longer finds himself in a reasonable position to find out that buying the shoes is morally impermissible. Maybe the shoe company has managed to temporarily purge all information regarding the matter from the internet. In missing the postings when they first appeared, Terry has squandered any opportunity to find out about the facts that could make his purchase deserving of moral blame. In this case, it may seem that, in unsubscribing, Terry has actually succeeded in putting himself in a position where he can no longer be blamed for doing something that he otherwise would have been blameworthy for. This seems untenable, however. Terry should not be able to deliberately avoid justified moral censure in this way. What explanation, then, can be given for why Terry is still to blame for buying the shoes?

4) The Regress at the Core of the Shirker Problem

The core of the shirker problem lies in this difficulty in pinpointing where exactly to locate the failure of attempted evasions of obligation like Terry's in UNSUBSCRIBE. Recall the Epistemic condition that we started with:

Epistemic condition (EC): For all putative obligations x , S is blameworthy for violating x only if S is in a position to know (with reasonable effort) that they are

¹³ I here ignore potential worries about the causal inefficacy of consumer choices that might undermine individual obligations not to purchase items produced under morally unacceptable conditions. For an overview, see Nefsky 2019.

subject to x.¹⁴

Given only this principle, it seems that Terry is not blameworthy for the act of buying the shoes. If EC and the obligation not to buy child-labour shoes were all we could go on, it would seem to follow that Terry actually ends up blameless for his purchase. If we want to hold onto both EC and the intuitive judgment that Terry is blameworthy, we must find some other obligation that he violates *earlier*, to which we can then later trace our ascriptions of blame for the act of buying the shoes.

The most straightforward solution for such an independent principle that Terry violates is first hinted at by Sorensen (1995: 255), endorsed by Ted Sider (1995: 278), and has been dubbed ‘Block’ by Wieland:

Block. I should refrain from making my obligation to refrain from buying the shoes unknowable.¹⁵ (Wieland, 2015: 291)

The worry, initially voiced by Sorensen and shared by Wieland, is that this is not a good solution to the problem as it generates an unfortunate regress. For, as an obligation, surely Block itself must be subject to EC. This opens up a loophole to be utilised by the shirker: As failing to meet Block itself is only blameworthy when Block is knowable to the agent, the obvious strategy for the shirker to follow is to try and make Block itself unknowable.¹⁶ Unfortunately, it would also not help if there were a further obligation in place (Wieland calls it SuperBlock) calling on us to keep Block knowable. After all, this obligation could in turn be evaded by shirking. In trying to explain the impermissibility of shirking through principles like Block, we thus inevitably seem to fall into an infinite regress. Since every such obligation is in principle shirkable, we end up with infinite chain of obligations not to shirk in our attempt to successfully ascribe a blameworthy benighting act to shirkers like Terry. If there is no obligation in the long chain that cannot be shirked, then shirking seems like a winning strategy. If this were true, we would likely want to reconsider subscribing to the Epistemic Condition in the first place. Indeed, this is the conclusion that Sorensen draws. For Sorensen, we are simply subject to many unknowable obligations and can be required to comply with them even when we are not in a position

¹⁴ Wieland frames his discussion with reference to a slightly different principle:

Access: For all obligations x, an agent s has x only if s can know she has x. (Wieland 2015: 289)

Although I prefer EC over Access both for reasons of generality and extensional adequacy, the differences between EC and Access are not relevant for the central argumentative moves in the following.

¹⁵ To specify: I think Block should be read as elliptical, really demanding to refrain from *deliberately* making the obligation unknowable. Otherwise, it would preclude many actions that we would not usually rate as blameworthy (I will return to this point later).

¹⁶ There are further complications that arise regarding Block, at least on the simple formulation laid out here. As spelled out here, Block is specifically concerned with the ‘target obligation’ of not buying the shoes, which crucially features in its content. That being so, making the target obligation unknowable should, ipso facto, make Block unknowable. In shirking on the obligation to buy the shoes, we would therefore automatically shirk on Block as well. I will ignore these further complications in what follows.

to know that we have them. If there are reasons not to blame agents for acting from a position of innocent ignorance, these are *indirect*: it might, for example, lead to better consequences not to blame innocently ignorant agents (Sorensen 1995: 269-71).

However, there are less radical suggestions for solving this problem which do not do away with EC. In the following section, I will present Wieland's solution and lay out why I find it unconvincing. Finally, I will present my own solution to the Shirker Problem.

Before doing so, let me quickly discuss a 'solution' that I think will not do, namely dismissing the problem by rejecting the regress as practically infeasible. It is true that no agent will likely ever find themselves in a situation where they have the opportunity to first decide to shirk on a concrete obligation, then decide to shirk on Block, then on SuperBlock, then on SuperSuperBlock, etc. However, as Wieland points out, it may well be that an agent can find themselves a position to shirk on all of these higher-order obligations at once, for example by avoiding any kind of engagement with moral philosophy, specifically those branches of it relevant to questions of blameworthiness under uncertainty (Wieland 2015: 298). And even if this turned out not to be practically feasible, it seems to me that the problem retains some bite even when removed from all possibility of ever actualising in the real world. What we are evaluating here is a way of thinking about blameworthiness, encapsulated in EC and a number of very plausible assumptions about what it means to be subject to moral obligations and how derivative responsibility functions. Showing the vicious regress at the heart of the Shirker Problem to even only be a *theoretical* possibility under these assumptions strongly indicates that these very assumptions need modifying or giving up. For this reason, we would do well to see if we can find an answer to the problem while retaining them, as I shall now attempt to do.

5) Self-Referential Obligations to the Rescue?

While he thinks that Block cannot save us from the Shirker Problem, Wieland continues the search for an obligation that can function as a solution to the problem. On his view, such an obligation would be required to have some very specific features. This has to do with how, according to Wieland, we can best make sense of what makes the regress at the heart of the Shirker problem troublesome. On Wieland's View, if we want to take issue with the regress and not simply regard it as benign, we must embrace what he calls 'Failure Theory' (Wieland 2015: 295). According to this theory, 'regress arguments demonstrate that solutions to a given problem fail because they get stuck in a regress, namely of similar problems that must be solved in order to solve the initial problem' (ibid.).

Based on this analysis, Wieland turns to searching for an obligation that can solve all

problems at once to evade the regress. He comes up with the following set of two obligations, one of which is importantly self-referential¹⁷:

Moderate General Law: O1: S should refrain from buying the shoes; O2: S should refrain from making or keeping O1 or O2 unknowable. (Wieland 2015: 297)

In Moderate General Law, O2 plays a dual function. For one, it forbids shirking on the particular obligation – in our case, the buying of the shoes. For another, however, the reflexive reference to the anti-shirking obligation O2 itself is meant to short-circuit the regress, eliminating the possibility of creating an excuse through deliberate abuse of the Epistemic Condition. This solution might seem surprisingly ad-hoc to some readers. This impression, however, might be alleviated when the problem is reconsidered in view of the technical intricacies of regress problems that heavily inform Wieland's discussion – here, solutions that do not have a flavour of ad-hoc-ness to them seem genuinely hard to come by. Ad hoc or not, I believe the technical solution employed by Wieland does not manage to provide a satisfactory solution to the Shirker Problem.

The central problem is one Wieland himself mentions (Wieland 2015: 297-98), as it has already been pointed out in the original discussion by Sorensen:

Can the infinite regress be short-circuited by interpreting the curiosity imperative as self-referential? No, because 'Everyone has a duty to learn his duties – including this very duty' is not self-evident. To infer it from itself would be circular. And to infer it from another proposition reopens the loophole; a shirker can avoid knowledge that the curiosity principle is true by disposing of his opportunities to learn whether it is true. (Sorensen, 1995: 255)

The central problem is the status of O2 itself. Is it justified? And much more importantly: is it accessible? It is important to note that O2 will only offer a solution to the problem if it is accessible to the shirker from the start. Its self-reflexive formulation may mean that it is ultimately not shirkable, but this does not mean it is automatically knowable by every agent.

In response to this point, Wieland paints a picture in which we are in a situation akin to a court of law in which the shirker is the accused, and we are prosecutor, judge, and lawmaker in unison:

Further suppose that I am the prosecutor or judge who wants to punish you, and finally that obligations are like laws that one has to institute, and that in the initial situation there are just two laws: Access [dialectically equivalent to EC, DBr] and Law A: Everyone ought to refrain from supporting child labour. (Wieland, 2015:

¹⁷ The term Moderate General Law is mine. Wieland does not himself label the set of O1 and O2 but reaches them after rejecting Weak General Law and Strong General Law as too weak and too strong, respectively.

In his opinion, the very fact that instituting Moderate General Law allows us to ‘nail’ the accused justifies its status as an obligation.

I am inclined to think that O2 is justified simply because it is part of a good solution to the Shirker Problem. Consider the law court analogy again: I am the judge, you the shirker; I have instituted O2 and now want to use it against you. However, now suppose you respond that you are not yet in the position to know O2, as I did not yet justify it, that is, did not yet convince you that it is a good law. Clearly, this response is inappropriate. All I need to secure is that you are in the position to know (and to be justified in believing) that you have to comply with O2, not why you must do so. (Wieland, 2015: 298)

This line of reasoning seems spurious for two reasons. First, I strongly doubt that the court of law analogy holds as a metaphor for justification of moral obligations. It is true that we are in a situation with a problem needing a solution. The idea that agents can excuse their morally wrong actions through epistemic shirking runs sharply contrary to our intuitions, while EC has a strong intuitive appeal. But that there is a problem to be solved here does not mean that any obligation we can come up with as a solution is automatically justified without any need for external support or further explanation.

Secondly, even if the court of law analogy passes, and O2 is justified simply because it is a good solution to the Shirker Problem, I do not see how this should make any difference for its *knowability*. And recall that knowability, in turn, is a prerequisite for it even being a solution at all. Of course, if we are in the actual situation of a court of law, and O2 is literally read out to a shirker, then this shirker will indeed be deprived of any possibility of legitimately shirking on their obligations. However, surely this cannot be the scope of solution that Wieland will content himself with. We do not only want a secure footing for blame in cases in which somebody attempts to shirk on an obligation and an ethicist then points out their mistake by reference to O2. We want to find out why shirking on obligations is *always* blameworthy, even if the shirker is never actually confronted.

I do not see how any of Wieland's arguments provide a case for the actual knowability of O2. It cannot be the case that being put forward as a solution to the Shirker Problem in the hypothetical court of law somehow makes O2 knowable in the real world. And this point generalizes. For any potential source of knowledge of O2 (such as a book of laws), defenders of moderately self-referential laws would have to show that this source of knowledge is necessarily available to all potential shirkers.¹⁸ The problem with Wieland's arguments is that neither they, nor anything in their vicinity, appear well-suited to furnish any explanation of the accessibility of O2. In the quotation above, Wieland is correct in

¹⁸ I thank an anonymous referee for Inquiry for pressing me on this point.

pointing out that not having independent justification does not mean the obligation is not knowable. But the reverse also holds. Knowability and justification are simply independent of one another. When Sorensen points out that his version of a self-referential obligation is not self-evident, he is making a point about knowability, not (only) justification.

At this point, it is important to keep in mind that the dynamics of the Shirker Problem clearly saddle the burden of proof on the defender of the universal principle such as Moderate General Law. Recall that we are trying to find a principle that someone explicitly concerned with avoiding morally relevant knowledge cannot help but have access to, no matter their efforts in curtailing their epistemic situation.¹⁹ What is more, there are some positive reasons for doubting that this burden of proof could be successfully discharged. Making a case that a general principle such as Moderate General Law is always knowable *ab initio* is a considerable task, given its complex content and very context-specific field of application. Even in philosophical circles, it is an obligation that I believe very few of those not familiar with the very narrow debate on the Shirker Problem would come up with easily. I thus conclude that, in absence of further independent argument for its universal knowability, Moderate General Law is not a good solution to the Shirker Problem, because O2 is not guaranteed to be accessible to the great majority of people.

Can these objections be avoided by adopting a stronger version of a self-referential obligation? If the problem is that we might not know our obligations from the start, are we possibly required to *learn* them? Consider the following strong self-referential obligation:

Strong General Law. S should *learn* all of her obligations (including this obligation). (Wieland 2015: 296)

I do not believe this obligation stands any better chance of solving the problem than Wieland's Moderate General Law. Like O2, Strong General Law is in need of independent justificatory support. What is more, Strong General Law gives rise to serious worries about extensional adequacy, as it appears to be implausibly demanding. At an extreme, it may it be contended that Strong General Law demands more than is humanly possible, and therefore straightforwardly conflicts with the independently plausible posit of Ought Implies Can. But even if the demands of Strong General Law are such that they could in theory be met, doing so certainly seems to be more than morality can reasonably require. As Sorensen quips, it would be a stretch to suggest that we have an obligation that would

¹⁹ Note that we thus not only need to show that the principle could be understood by the shirker (if it was read out, for example), but that the shirker could come to know this principle on their own, notwithstanding their persistent attempts of epistemic self-curtailment.

prescribe extensive ‘mandatory ethics courses’ to everyone (Sorensen 1995: 255).

And an even more serious worry looms for the proposal under consideration. Even if these first two problems can be circumvented, it is not clear how Strong General Law could even be marshalled in an attempt to save EC in the first place. If we add the proviso that failing Strong General Law can be blameworthy, the principle rather seems to simply stand contrary to EC. It is an obligation that, by its very nature, applies to agents who are not (yet) in a position to know it, the point of which is to find grounds for blame against these unwitting agents. But this is the very kind of blame EC is supposed to rule out. In going down this route, we therefore seem to be at high risk of throwing out the baby with the bathwater.

For these reasons, I thus suggest that we should content ourselves with the weaker claim: we have an obligation not to consciously make our obligations unknowable in Shirker-like cases. However, as I have also argued, simply making this obligation self-referential won't solve the Shirker Problem for us. In the following, I will show why we ultimately do not need any such exotic remedy. The solution to the Shirker Problem, I suggest, turns out to lie in the very idea of what it means to epistemically shirk on an obligation.

6) Motivations for Shirking and Grounds for Justified Blame

To find a solution to the Shirker Problem, we must first return to closely examining a given agent's reasons and motivations for their decision to shirk. Since the kinds of wilful ignorance that can truly be called shirking are ones in which ignorance is caused in reaction to a suspicion about (future) moral obligations, a shirker always acts in conscious infringement on a moral obligation, or at least conscious embrace of *risk* of infringing on it.

To show why this is the case, let me first introduce some useful taxonomy by Rik Peels, which can help shed light on this matter.²⁰ Peels distinguishes between different kinds of propositional ignorance, four of which will be relevant to the discussion in question: *Disbelieving Ignorance* occurs iff (i) it is true that p , and (ii) an agent S disbelieves that p . *Suspending Ignorance* iff (i) it is true that p , and (ii) S is unsure whether or not p , and therefore suspends belief. *Undecided Ignorance* occurs iff one has considered p but fails to adopt a doxastic attitude toward p (due to distraction or any other interruption of the process of deliberating on whether or not p). Finally, we can talk about *Deep Ignorance* in cases which the question whether or not p never occurred to the agent.

²⁰ Peels laid out this taxonomy in an initial form in Peels 2014 and then in a revised version in Chapter 4 of Peels 2023.

Formally, these would be cases in which (i) it is true that p , and (ii) S neither believes that p , nor disbelieves that p , nor suspends belief on p .²¹

In brief, my argument is as follows: A) There is an important class of morally relevant propositions that epistemic shirkers can only be suspendingly or undecidedly ignorant of. B) Refraining from taking into account one's suspending or undecided ignorance regarding morally relevant matters in one's decision-making amounts to acting in a way that is morally risky. C) Anyone behaving in such a morally risky manner can justifiably be blamed for doing so. Therefore, any genuine decision to epistemically shirk on an obligation is necessarily blameworthy. In what follows, I will defend the premises of this argument in turn.

A) *Shirkers are subject to suspending or undecided ignorance at best.*

Let us begin with the claim that shirkers necessarily display only suspending or undecided ignorance regarding some crucial morally relevant propositions. This claim straightforwardly falls out of any reasonable definition of shirking. Recall that shirkers are agents who *deliberately* deteriorate their epistemic position to evade justified blame. In order for an attempt of effecting a negative change to one's future epistemic state being properly called an instance of epistemic shirking, the agent thus cannot be deeply ignorant of the possibility of a future infringement of a moral obligation. To evade knowledge about a particular fact of the matter F *because* of its implication for one's future moral obligations, it must have at least occurred to one that there *might* be certain moral obligations that depend on this knowledge.

Similarly, it would seem strange to call someone an epistemic shirker if they are disbelievingly ignorant about any relevant obligations from the getgo, i.e. of the firm and honest conviction that there are no moral obligations that depend on the future knowledge they cause themselves not to have.²² In Wieland's original example of epistemic shirking, a person fails to watch the 8-o'clock news, and then unknowingly buys shoes produced

²¹ In Peels 2023, Peels further distinguishes three subforms of what he originally termed "deep ignorance": *Unconsidered Ignorance*, *(True) Deep Ignorance*, and *Complete Ignorance* (Peels 2023: 78-82). Albeit interesting in their own right, these distinctions do not make a difference for the arguments to follow, which is why I will continue using the term "deep ignorance" as an umbrella term encompassing all three subforms, as in Peels' original 2014 contribution.

²² That is not to say that shirkers could not be disbelievingly ignorant of any proposition, which is why it is important to qualify my statement here. My claim is that epistemic shirkers are necessarily only suspendingly or undecidedly ignorant of propositions such as "I could learn about *some* important obligation by pursuing course of action A". When it comes to concrete statements of which particular obligations x , y or z these may be, shirkers may very well be blamelessly disbelievingly ignorant. However, their awareness of the general possibility of learning more, coupled with the likelihood that this will overturn their evidential situation with respect to one or more of these more concrete propositions, grounds a requirement not to actively avoid exposure to further evidence. Thanks to an anonymous referee for Inquiry for pressing me to clarify this matter.

in dubious circumstances because she missed some information about working conditions imparted by the programme. Assuming the agent in the example had no special reason to expect morally crucial information in the news broadcast, it seems to me that this agent would be fully excused for buying the shoes.

Crucially, this could even be the case if we explicitly construct this case as one of wilful ignorance, i.e. ignorance deliberately chosen by the agent. Take, for example, Tammy, who wants to keep her mind ‘clean and unburdened’ before an important exam on the following day. She thus decides to take in as little new knowledge as possible on the night leading up to the exam, even if this means breaking her pleasant daily habit of watching the 8-o’-clock news. As long as Tammy is (innocently) deeply or disbelievably ignorant that the knowledge she is avoiding has significant moral relevance, we would not criticise her actions as an instance of shirking, even if the newscast she avoids contains the very same information as Terry’s celebrity post. Of course, if we subscribed to something like Strong General Law, we could have grounds for also holding agents like Tammy responsible. But not only would this strongly go against our settled intuitions in this case, but furthermore saddle us with substantial philosophical costs, as I have already laid out above.

Between the two extreme examples of Terry’s akratic decision to impermissibly defy a risk and Tammy’s innocent avoidance of information she expects to be of low moral relevance, a wide range of cases of course remains.²³ Not every shirker will be subject to clear-eyed akrasia when engaging in shirking behaviour. The shirking decision to avoid the potentially cumbersome information may be born from inattentiveness, laziness, or other kinds of epistemic or practical vice. Crucially, however, even those shirkers who manoeuvre themselves into bad epistemic positions as a result of some kind of negligence will necessarily have been in a position to recognize their behaviour as indeed negligent. To count as genuine shirkers, they must have considered the potential gains of morally relevant knowledge that they deliberately avoid.²⁴

On the other hand, there may be cases where the moral evaluation of agents engaged in shirker-like behaviour is not as clear-cut. As I mentioned at the beginning, a convincing version of EC must include a reasonability clause to allow for the blamelessness of agents who would have to face excessive costs to gain some morally relevant knowledge. Some shirkers may make a wrong call with respect to the evaluation of these costs, yet do so in understandable ways. Imagine a cross between Terry and Tammy who unsubscribes from

²³ I thank two anonymous referees for Inquiry for pressing me to consider these intermediate cases.

²⁴ As laid out in section 2 above, providing a full account of what it takes to be in a position to know a given proposition unfortunately goes beyond the scope of this paper. However, it appears to me that any precisification of the relevant notion of epistemic accessibility worth its salt should yield the accessibility claim at the heart of my argument: Shirkers, as agents who must have considered that there may be morally relevant information at stake in *p*, are in a position to know that avoiding information pertinent to *p* carries a moral risk.

the feed not eagerly, in order to avoid ethical challenges to their lifestyle, but reluctantly because they wish to avoid informational distraction that could hamper their professional achievements. Depending on details, such an agent may be at least partially excused. Such intermediate cases can be complex and I regret not being able to engage with them in more depth here. However, I take it that the paradigmatic case of shirking, and certainly the one that is at issue in the debate between Sider, Sorensen and Wieland, is one that bears closer resemblance to Terry's case. And here, the situation is clear.

To sum up: it is necessary for someone to be rightfully called a shirker that they be aware of their having some (possibly future) obligation x, or at the most having considered this matter and not come to the honest conclusion that there is no such obligation. Shirking is a *reaction* to the possibility of x, or something like it, occurring to them.

B) Acting from suspending ignorance is risky.

Grounds for blameworthiness can be found in each instance of epistemic shirking precisely because shirkers are at best suspendingly or undecidedly ignorant of the fact that their actions stand to affect their compliance with future obligations. For, as I shall argue now, suspending or undecided ignorance, unlike deep or disbelieving ignorance, does not usually excuse agents from wrongdoing.

This point is perhaps most quickly established by example. On New Year's Eve in Berlin, it is an unfortunately common practice amongst some people to throw lit fireworks into subway entrances for laughs. It is not hard to grasp why this is a morally problematic past-time: in engaging in this activity, one stands to cause serious harm to innocent passers-by in the subway tunnel. At the moment of throwing, however, perpetrators generally *do not know* whether or not their projectile is going to harm somebody. We can assume that they do not have the necessary evidence to firmly believe the sentence 'By throwing the firework, I will harm someone'. They are thus ignorant of the fact whether or not the action they are about to commit has morally problematic consequences. But this particular ignorance, real as it may be, surely cannot excuse their behaviour. In cases of suspending or undecided ignorance like these, choices are subject to considerations of risk. You are morally permitted to throw fireworks into a subway entrance (for whatever reason) only if you are *sufficiently sure* that there is nobody there – usually, because you checked thoroughly. Suspending or undecided ignorance about situational features does not furnish agents with an excuse for breaching obligations if what is at stake is of

considerable moral importance. Only deep or disbelieving ignorance can really render the agent blameless.²⁵

Let us return to our shirker. As I see it, the situation for them is as follows: they are suspendingly or undecidedly ignorant about (or even aware of) the existence of some future obligation *x*, and now must decide whether or not to avoid knowledge of some fact *p* that is relevant to this obligation. Course of action A, in which they do not shirk, will lead them to learn *p* and thus enable them to later act on that knowledge and discharge *x*, an obligation that at the moment of the decision they take themselves at least possibly to have. Course of action B, in which they do shirk, will lead them to not acquire knowledge of *p*, and thus prevent them from later forming the intention to do right by acting according to *x*.²⁶

C) Defiance of risk is blameworthy.

At the moment of the decision, the shirker thus clearly is in a position to know that choosing B might mean that they will act badly in the future, though of course they believe they will not be blameworthy for it. With A, on the other hand, there is a morally acceptable alternative available to them that will actually cost them *less* effort – recall that it is in choosing B that they have to make an effort to actively avoid knowledge of *p*. The fact that B is in this sense necessarily more effortful should also help alleviate worries that at least in some instances of epistemic shirking, self-regarding reasons can legitimately outweigh the moral reasons not to shirk. Of course, regarding future self-regarding reasons, things might look very different: course of action A, which leads the agent to recognise potentially demanding moral duties, might lead to more future effort for the agent, either with regards to compliance with these duties or with regards to dealing with guilt and perhaps blame. But if self-regarding reasons concerning this latter effort would outweigh any reason to act according to obligation *x*, it would be wrong to say that the agent is all-things-considered required to act according to *x* anyway, meaning that shirking would be redundant in this case. Again, the necessary conditions for an act to be accurately described as shirking ensure the blameworthiness of this act.

²⁵ For others coming to similar conclusions, see Peels 2014: 13; 2023, Chapter 10 and Harman 2011: 447-450. Alex Guerrero (2007) has argued that there is no reason why this important insight cannot be transferred to moral knowledge as well. Based on such considerations of risk, he concludes a principle of ‘Don’t Know, Don’t Kill’, which forbids taking the lives of beings when one cannot with sufficient confidence conclude that they do not have moral status.

²⁶ We have to be a little careful here. It is not the case that following course B makes it impossible for the shirker to fulfil the obligation *x* (see footnote 10). Of course, it does make the fulfilment of the obligation a lot less likely. What taking course B *does* make impossible is to do as *x* demands *because of its moral import*.

Our shirker thus clearly *must* find themselves in an epistemic situation that allows them to conclude that they are making a bad moral decision. The fact that this knowledge is accessible to them follows from the very definition of what it is to be a shirker. Whether or not their self-effected ignorance would at the later time be such that they would count as excused according to EC, we can therefore confidently trace back to this disregard of morally salient considerations as a blameworthy *benighting act*. This will ensure derivative responsibility for their later failing to act in a way as the obligation demands or would have demanded.

Some might now object that there is an important difference between epistemic shirking situations like UNSUBSCRIBE and ordinary risk-based cases like the fireworks example. While the person throwing fireworks into the subway is aware of the obligation they are possibly disregarding, Terry, when considering whether or not to unsubscribe, might have no idea exactly what kind of obligation he would discover were he to continue following the singer. However, I do not think that knowing what kind of moral obligation one will infringe upon makes a difference to the question of whether we are blameworthy for our choices. Imagine that a crazy pharmacist offers you a drug that promises 2 hours of a very pleasant high. However, the drug has weird and unsavoury side-effects. In one third of the cases, it will lead you to break a promise you made to someone. In one third, you will be caused to unwittingly punch someone innocent in the face, and in the final 33.33%, you will blurt out a racist slur in public. Not knowing which obligation you will violate clearly does not make a difference to the moral status of your decision of whether or not to take the pill. You are just as firmly required not to take the lottery pill as you would be required not to take a pill that will cause you to punch innocents in 100% of cases. What matters is that you risk doing wrong. The fact that there are moral obligations at issue in your choice, whatever they turn out to be, is enough to know that you should choose an alternative course of action.

As it turns out then, discussing the matter of epistemic shirking in terms of ignorance of a *certain* obligation, as well as ‘knowability’ of this obligation, can obscure the differences between different kinds of propositional ignorance. It is true that at the moment of the decision to shirk, before the possible acquisition of knowledge that p, obligation x is strictly speaking not yet fully knowable to the agent. As we have seen, however, in cases of suspending or undecided ignorance and risk of moral wrongdoing, full knowability is not a necessary prerequisite for blame. In all cases of shirking, the shirker is neither disbelievably nor deeply ignorant of the fact that he is engaging in behaviour that risks culminating in his infringing on *some* moral obligation. The very definition of shirking guarantees that this fact is accessible to them.

One final worry that might be brought up here is that this fact is unsuitable as a motivating reason for the agent to act on because it displays an unfortunate kind of moral

fetishism. All that the agent deliberating on the question of whether to shirk can go on is the fact that choosing A will enable them to later choose doing ‘the right thing’, whatever that turns out to be. However, it may be argued, this is not what good people usually care about. As Michael Smith (1994: 76-78) has famously argued, good people care about doing what is right *de re* (i.e., doing the things that help people, that respect people’s dignity, that avoid suffering etc.), and not just *de dicto* (i.e. whatever action is correctly labelled as ‘morally right’).

This worry does not strike me as decisive, however. For one thing, whether there is really something wrong about being motivated by what is right *de dicto* is itself a matter of heavy contention.²⁷ For another, it seems not out of the question that the situation of the shirker does not actually allow for a *de re* motivation. In the given situation, the potential shirker could for example imagine multiple possible scenarios and then act upon the moral considerations relevant in some of them – even if the probability that the envisioned scenario will eventually materialise is ultimately quite low. In fact, the cases discussed in the foregoing show that we simply must reject any requirement of *de re* motivation that is strong enough to exclude choices to not put oneself in a situation where one expects to act immorally in unexpected ways – such a view would simply yield unconscionable recommendations in many of the cases we considered.

To wrap up, let me draw together the strands just outlined again, which allows me to specify exactly how we can justify our reactions of blame to shirking agents like Terry. Even though at t_1 Terry is not in a position to know that he should not buy the shoes, Terry *is* in a position to know that he is risking disregarding morally relevant considerations – if he was not in this position, we could not justifiably consider him an epistemic shirker. At best, he would be an epistemically unlucky agent accidentally stunting their morally relevant knowledge like Tammy. Since he is in a position to know there is a risk, and since risking disregarding morally relevant considerations is wrong, he is likewise in a position to know he is violating his obligation not to incur such risks, which is not identical to any obligation not to buy the shoes. We can thus trace back blame for Terry’s failure to avoid supporting slave-labour through his purchase to this earlier failure of not showing proper regard for moral considerations more generally.

What, however, of this latter fact, the fact that it is wrong for Terry to risk doing moral wrong, whatever it may be – is this not something ignorance of which could again get Terry of the hook? No, it is not. For one thing, it is not a principle which Terry could successfully shirk – if Terry decides to stop himself from knowing more about the moral relevance of risk on a suspicion that risk might somehow be morally relevant, Terry must already believe there to be at least potential moral connect here – he is thus necessarily

²⁷ See for example Lerner 2018, Johnson King 2019 and Hering 2024.

already in a position to know that accepting a risk of one's own future disregard of morally relevant considerations could itself be morally wrong. But what if Terry is already deeply or disbelievingly ignorant of any requirement to refrain from running moral risks? In this case, I think we really might have to let Terry off the hook for shirking. It is, however, important to note that such deep or disbelieving ignorance is the mark of a character that is so extremely far removed from the ordinary to be close to bordering on moral insanity. Someone who firmly believes that it is wrong to murder someone but concurrently believes that there is absolutely nothing speaking against putting oneself in a position where there is a 50% likelihood that one will murder someone is likely so far removed from the standards of our normal moral community as to be exempt from any blame on account of this very fact.²⁸

7) Conclusion

Must we worry about epistemic shirkers? If what I have argued above is correct, then the answer is no, at least when limiting ourselves to theoretical worries that are the focus of the Shirker Problem. The reason for this is not to be found in complicated and demanding anti-shirking-principles with universal scope and potentially self-reflexive content, but instead in the nature of the motivations necessary for genuine epistemic shirking and the moral relevance of risk. Since every genuine shirker finds themselves in a situation where it must have occurred to them that the possibility of their shirking behaviour could lead to an infringement of moral obligation, we can confidently trace back blame to their objectionably reckless decision to shirk.

Admittedly, this need not be the only way in which they fall short of the normative mark. Many shirkers will already have access to the knowledge necessary for us to successfully apply norms like Block or more universal anti-shirking-laws to them (as noted above, genuine cases of fully pre-emptive shirking are likely to be quite rare). What is more, agents engaging in any form of wilful ignorance will often display forms of vice, both moral and epistemic, which can give us grounds for independent criticism. It is in the theoretically especially interesting instances where these further factors are absent that were the focus of discussion here.²⁹

That being said, attempts at forms of epistemic shirking remain an important, and likely growing, *practical* problem. As I have hinted at above, there is good reason to believe that attempts at conscious evasion of information pertinent to one's moral duties

²⁸ For a classic treatment of these kind of exemptions due to what we may call moral insanity or pathological evil, see Watson 2004.

²⁹ I thank an anonymous referee for Inquiry for pressing me to say more about alternate grounds for blameworthiness in many instances of shirking behaviour.

are widespread and contribute significantly to many pressing moral issues of our time. In response, we must find efficient political and social strategies to combat the social and technological trends facilitating such deliberate epistemic curtailment. This is not exclusively, and perhaps not even primarily, a job for moral philosophy. However, moral philosophy can contribute to efforts in curbing wilful ignorance by providing their theoretical foundations. If what I have argued above is correct, we need not to know substantive further details to come to a well-founded judgment for any case of deliberate epistemic shirking, even in the most improbable fully pre-emptive cases: When it comes to strictly moral blameworthiness, epistemic shirking will simply not get you off the hook.³⁰

References

- Andrić, Vuko. 2021. *From Value to Rightness: Consequentialism, Action-Guidance, and the Perspective-Dependence of Moral Duties*. London: Routledge.
- Calhoun, Cheshire. 1989. "Responsibility and Reproach." *Ethics* 99 (2): 389–406. doi:10.1086/293071.
- Code, Lorraine. 2014. "Culpable Ignorance?" *Hypatia* 29 (3): 670–76. doi:10.1111/hypa.12071.
- Fischer, John M., and Neal A. Tognazzini. 2009. "The Truth About Tracing." *Noûs* 43 (3): 531–56. doi:10.1111/j.1468-0068.2009.00717.x.
- Guerrero, Alexander A. 2007. "Don't Know, Don't Kill: Moral Ignorance, Culpability, and Caution." *Philosophical Studies* 136 (1): 59–97. doi:10.1007/s11098-007-9143-7.
- Harman, Elizabeth. 2011. "Does Moral Ignorance Exculpate?" *Ratio* 24 (4): 443–68. doi:10.1111/j.1467-9329.2011.00511.x.
- Harman, Elizabeth. 2015. "The Irrelevance of Moral Uncertainty." In *Oxford Studies in Metaethics Vol. 10*, edited by Russ Shafer-Landau. First edition. Oxford: Oxford University Press.
- Heering, David. 2024. "Why and When Is Pure Moral Motivation Defective." *Erkenntnis* 89 (2): 665–684. doi: 10.1007/s10670-022-00550-w
- Johnson King, Zoë. 2020. "Praiseworthy Motivations." *Noûs* 54 (2): 408–30. doi:10.1111/nous.12276.
- Kiesewetter, Benjamin. 2017. *The Normativity of Rationality*. First edition 1. Oxford: Oxford University Press.
- Kiesewetter, Benjamin, and Jan Gertken. 2021. "How Do Reasons Transmit to Non-Necessary Means?" *Australasian Journal of Philosophy* 99 (2): 271–85. doi:10.1080/00048402.2020.1745252.

³⁰ This paper has had a comparatively long history. Over the years, I have received helpful comments from many people on drafts in wildly different stages of development. I thank Maike Albertzart, Francesca Bunkenborg, David Heering, Benjamin Kiesewetter, Thomas Schmidt, Razvan Sofroni and two anonymous referees for Inquiry for their helpful written comments. Material from this paper was also presented at conferences and colloquia in Berlin and Gothenburg. I am grateful to audiences and organisers for the helpful exchanges and feedback. Work on this paper was funded by Deutsche Forschungsgemeinschaft (Project number: 503995582).

- Lerner, Adam. 2018. "The Puzzle of Pure Moral Motivation." In *Oxford Studies in Metaethics Vol. 13*, edited by Russ Shafer-Landau. Oxford: Oxford University Press.
- Lord, Errol. 2015. "Acting for the Right Reasons, Abilities, and Obligation." In *Oxford Studies in Metaethics Vol. 10*, edited by Russ Shafer-Landau. Oxford: Oxford University Press.
- Lynch, Kevin. 2016. "Willful Ignorance and Self-Deception." *Philosophical Studies* 173 (2): 505–23. doi:10.1007/s11098-015-0504-3.
- Martín, Annette. 2021. "What Is White Ignorance?" *The Philosophical Quarterly* 71 (4). doi:10.1093/pq/pqaa073.
- Medina, José. 2013. "Epistemic Responsibility and Culpable Ignorance." In *the Epistemology of Resistance: Gender and Racial Oppression, Epistemic Injustice, and Resistant Imaginations*. Studies in feminist philosophy. New York: Oxford University Press.
- Mills, Charles W., ed. 2017. *Black Rights/White Wrongs*. Oxford: Oxford University Press.
- Moody-Adams, Michele M. 1994. "Culture, Responsibility, and Affected Ignorance." *Ethics* 104 (2): 291–309.
- Nefsky, Julia. 2019. "Collective Harm and the Inefficacy Problem." *Philosophy Compass* 14 (4): e12587. doi:10.1111/phc3.12587.
- Peels, R. 2014. "What Kind of Ignorance Excuses? Two Neglected Issues." *The Philosophical Quarterly* 64 (256): 478–96. doi:10.1093/pq/pqu013.
- Peels, Rik. 2023. *Ignorance: A Philosophical Study*. Oxford: Oxford University Press.
- Sarch, Alexander. 2018. "Willful Ignorance in Law and Morality." *Philosophy Compass* 13 (5): e12490. doi:10.1111/phc3.12490.
- Sider, Theodore. 1995. "Sorensen on Unknowable Obligations." *Utilitas* 7 (2): 273–79. doi:10.1017/s0953820800002089.
- Simon, William H. 2005. "Wrongs of Ignorance and Ambiguity: Lawyer Responsibility for Collective Misconduct." *Yale Journal on Regulation* 22 (1): 1–35.
- Smith, Holly. 1983. "Culpable Ignorance." *The Philosophical Review* 92 (4): 543–71. doi:10.2307/2184880.
- Smith, Michael. 1994. *The Moral Problem*. Oxford: Blackwell.
- Sorensen, Roy. 1995. "Unknowable Obligations." *Utilitas* 7 (2): 247–71. doi:10.1017/S0953820800002077.
- Watson, Gary. 2004. "Responsibility and the Limits of Evil: Variations on a Strawsonian Theme." In *Agency and Answerability: Selected Essays*. Oxford: Clarendon.
- Wieland, Jan W. 2015. "Access and the Shirker Problem." *American Philosophical Quarterly* 52 (3): 289–300. <http://www.jstor.org/stable/24475464>.
- Wieland, Jan W. 2017a. "Responsibility for Strategic Ignorance." *Synthese* 194 (11): 4477–97. doi:10.1007/s11229-016-1145-6.
- Wieland, Jan W. 2017b. "Willful Ignorance." *Ethical Theory and Moral Practice* 20 (1): 105–19. doi:10.1007/s10677-016-9722-9.
- Williams, Daniel. 2021. "Motivated Ignorance, Rationality, and Democratic Politics." *Synthese* 198 (8): 7807–27. doi:10.1007/s11229-020-02549-8.
- Williams, Nancy M. 2008. "Affected Ignorance and Animal Suffering: Why Our Failure to Debate Factory Farming Puts Us at Moral Risk." *Journal of Agricultural and Environmental Ethics* 21 (4): 371–84. doi:10.1007/s10806-008-9087-8.