

Anna Hartford\*

# How Much Should a Person Know? Moral Inquiry & Demandingness

<https://doi.org/10.1515/mopp-2018-0056>

**Abstract:** An area of consensus in debates about culpability for ignorance concerns the importance of an agent's epistemic situation, and the information available to them, in determining what they ought to know. On this understanding, given the excesses of our present epistemic situation, we are more culpable for our morally-relevant ignorance than ever. This verdict often seems appropriate at the level of individual cases, but I argue that it is over-demanding when considered at large. On the other hand, when we describe an obligation to know that avoids over-demandingness at large, it fails to be sufficiently demanding in individual cases. The first half of this paper is dedicated to setting up this dilemma. In the second half, I show that it cannot be easily escaped. Finally, I suggest that this dilemma impedes our ability to morally appraise one another's ignorance, and even our own.

**Keywords:** ignorance, demandingness, information excess, moral inquiry

## 1 Introduction

Ignorance is a morally confounding state: we often excuse people precisely because they 'didn't know' or they 'didn't realise,' but at other times we intend ignorance as censure, and consider it blameworthy in itself. The problem of ignorance, and our culpability for ignorance, has recently received significant philosophical attention.<sup>1</sup> But relatively little attention has been paid to one of the central concerns in our contemporary moral and epistemic lives: what are our obligations regarding the sheer *amount* of information we now have easy access to, much of it of potential moral relevance?<sup>2</sup>

---

<sup>1</sup> This attention culminated recently in the publication of Robichaud and Wieland (2017).

<sup>2</sup> For other treatments of this question, see Vanderheiden (2016), and Dennett (1986) (I will return to both).

---

\*Corresponding author: Anna Hartford, Department of Philosophy, Stellenbosch University, Private Bag X1, Matieland 7602, South Africa, E-mail: [annahartford@gmail.com](mailto:annahartford@gmail.com)

An area of consensus in debates about culpability for ignorance regards the importance of an agent's epistemic situation in determining what they ought to know. Where the relevant information is difficult to obtain, we are more easily excused in our ignorance; conversely, where the relevant information is easily accessible, blame is more appropriate. On this understanding, given the excesses of our present epistemic situation, we are more culpable for our morally-relevant ignorance than ever. This verdict often seems appropriate at the level of individual cases, but I will argue that is over-demanding when considered at large. On the other hand, when we describe an obligation to know that avoids over-demandingness at large, it fails to be sufficiently demanding in individual cases.

The first half of this paper is dedicated to setting up this dilemma. An important initial question concerns the scope of our obligations to know, which are often thought of as deriving from other moral obligations, and tracking their strength and seriousness. These considerations are relevant to the question of demandingness, since they seem to guide and limit our obligations to know. Where we are aware of our ignorance, and its potential moral repercussions, it seems clear that this knowledge ought to inform our efforts in inquiry. But what happens when we are *unaware* of the moral repercussions of our ignorance, or even unaware of our ignorance? I argue that our obligations to know must include obligations to uncover our ignorance (to the extent that we are able) and its potential consequences. This is the problem of 'deep ignorance,' or the 'unknown unknown,' and together with the nature of our epistemic situation, I argue that it generates an over-demanding obligation to know.

In the second half of this paper, I endeavour to show that the dilemma I have sketched cannot be easily escaped: I consider a variety of ways of limiting the realm of our responsibility to know within our epistemic context, but argue that they are ultimately unsatisfactory. Finally, I suggest that this dilemma impedes our ability to morally appraise one another's ignorance, and even our own. Where we are excused not by the dearth of easily-accessible information, but by its abundance, it becomes difficult (if not impossible) to determine whether a particular instance of ignorance should be deemed blameworthy or not.

## Section 2

### 2.1 Two paradigms concerning culpability for ignorance

The question of how much we should know can also be posed as the question of when we become culpable for what we do not know. Two central paradigms

have emerged concerning culpability for ignorance (and, arguably, for the wrongdoings we commit within ignorance). Loosely put: one concerns our actions, or what we do, and the other concerns our pattern of moral concern, or who we are.

To begin with the first paradigm, concerning action: we are able to alter, expand and complicate our beliefs by inquiring, researching, informing ourselves, reflecting, considering opposing views, and so on. Say Bob runs a factory that produces massive amounts of waste which he releases into the surrounding environment. It seems patent that Bob ought to investigate the effects of this waste on the neighbouring community. Should Bob fail to investigate, and it is subsequently discovered that the waste has led to many cases of serious illness, the excuse of ignorance would be impotent: Bob may not have known, but clearly he should have known.

But the relationship between inquiry and subsequent moral action is far from straightforward. After all: one could inquire, discover that some action is wrong, and proceed all the same. Or one could fail to inquire, and cause no harm purely by fluke.<sup>3</sup> One can even inquire and still come to hold an errant belief. Let's say Bob undertakes an investigation into the waste, which results in a near consensus that it is toxic and that in the quantities Bob is producing it will be devastatingly harmful. But the investigation also resulted in some inconsistencies, and inconclusive findings; it also seems that in small enough quantities the waste isn't harmful. Deliberating on the information revealed by his inquiry (and fixating especially on the few inconsistencies), Bob decides that most likely no harm will be caused, and that it is fine for his factory to proceed as before. Let me stipulate that Bob ultimately comes to this errant belief precisely because it justifies his present practices and will cost him less money. It seems odd to say, in this case, that Bob's ignorance is blameless merely because he undertook an inquiry. It is quite clear that despite his inquiry, Bob's ignorance is being maintained and sustained by self-interested motives and moral disregard that have implications for his blameworthiness. This leads us to a second sense in which we can be blameworthy for our ignorance, which turns not on our actions, but rather on our pattern of moral concern.<sup>4</sup> From this

---

<sup>3</sup> See Wieland, 2017a p. 8. Wieland has separately suggested that inquiry enables us to non-luckily comply with our moral obligations, in a way that has significance for assessments of praiseworthiness and blameworthiness. (It might not be the case that one can only non-luckily comply under these circumstance: if we consider inverse akrasia cases, we might think that agents have non-luckily complied, even though they are mistaken in their beliefs, because their right actions have emerged from the right-making features of their act. See Arpaly & Schroeder 1999).

<sup>4</sup> This is usually captured by *Quality of Will* accounts of blame.

perspective, we are blameworthy for our ignorance insofar as it manifests ill will, or insufficient regard. (Where someone acts out of malice, cruelty or spite this insufficiency is quite obvious, but it can also be present when someone acts out of indifference or negligence).

These two perspectives concerning culpability for ignorance come apart in important ways.<sup>5</sup> Yet they are also often related: if Bob failed to investigate, it seems quite clear that he would be demonstrating insufficient concern for whoever he might harm. Similarly, if Bob arrived at an errant belief despite his investigation, our conviction that his ignorance evinces insufficient concern is tied up with our conviction that something has gone amiss in his deliberations. It seems that having the right sort of moral concern, while separate from questions of inquiry, also involves having a certain quality of attention.<sup>6</sup> For my present purposes, the relationship between these two perspectives is more important than the differences, and I will draw on both of them in the conversation which follows. (As we will soon see, both assessments are also fundamentally responsive to an agent's epistemic situation, and both assessments therefore have implications for the parameters of our blameworthiness within our present epistemic situation).

Having introduced these paradigms, let me turn to questions of the scope and strength of our obligation to know. These considerations are important to the present discussion, since they could conceivably constrain the obligation, and rescue it from the charge of over-demandingness. I will proceed by arguing that the problem of deep ignorance, or the unknown unknown, seriously undercuts the success of these constraints.

## 2.2 The scope of our obligation to know

The realm of our knowledge will always be paltry compared to the realm of our ignorance, yet much of this ignorance seems morally benign. I know that there is a tree outside my window, but I couldn't tell you what kind of tree it is, or when it was planted, or how many leaves it has, or which birds prefer its fruits. I am ignorant on all these fronts, and yet none of this ignorance seems to matter much. In the boundlessness of all that we do not know, how ought we to determine the small realm of ignorance that we are *morally* at fault for?

---

<sup>5</sup> In particular these paradigms might have different implications for our blameworthiness for ignorance we have no conscious control over.

<sup>6</sup> I am following Vanderheiden (2016) in using the phrase "obligation/s to know." I intend this phrase to capture both paradigms concerning culpability for ignorance.

An initial suggestion is that the obligation to know is derivative: it derives from, or is generated by, our other moral obligations. Gideon Rosen, who in some senses initiated the contemporary debate about responsibility for ignorance, writes that: 'We are under an array of standing obligations to inform ourselves about matters relevant to the moral permissibility of our conduct: to look around, to reflect, to seek advice, and so on'.<sup>7</sup> So to return to the example from earlier: Bob's ignorance regarding the toxic waste is wrong precisely because harming people is wrong. On the other hand, Bob needn't know anything about the Mexican-American War, or the featheriness of dinosaurs, or the process of the Ricci flow. All he needs to know, in this case, is whether harm will result from his manufacturing.

The derivative construal drastically reduces what we are morally at fault for not knowing (even if we separately agree that much of this other information has value). It also seems to perform a great deal of explanatory work. For a start, it explains why we have stronger obligations to inform ourselves in some areas rather than others. It seems clear, for instance, that certain obligations to know arise with the assumption of professional duties: a doctor is obliged to keep up to date with the medical literature; whereas the rest of us have no such obligation. More broadly, we might think that obligations to know are generated from a role like citizenship, especially insofar as our electoral and other choices as citizens have moral implications. It also seems plausible that there might be proximal and temporal dimensions to our moral obligations to know, which is explained by the derivative construal insofar as our acts and choices have a greater moral impact proximately and, of course, presently (although both of these assumptions have limitations).

The derivative construal implies that we are morally obliged to inform ourselves when the outcome of this inquiry will have some sort of moral value, and we are not morally obliged to inform ourselves when it does not. It has also been suggested that obligations to know should track the seriousness of the moral issues which they concern. So where the moral seriousness of a decision is significant, the related obligation to know is similarly significant. (On the other hand where there are very low moral stakes, the obligation is weaker). Alexander Guerrero has advanced a position along these lines: Guerrero's primary claim is that 'what is required of us from an epistemic point of view increases as what is at stake from a moral point of view increases: our epistemic obligations grow as the relevant context becomes more morally serious'. (2007, p. 70)

---

<sup>7</sup> Rosen (2003, p. 65). Holly Smith (2014) has made a more detailed case for this derivative construal.

This is an intuitive notion, and it has been widely adopted. And where we are aware of the moral stakes of our ignorance, it seems patent that this awareness should guide our inquiry. But note, crucially, that we can be, and often are, *unaware* of the moral consequences of our ignorance, or even unaware of our ignorance at all. I will refer to this, following Guerrero (2007), as ‘deep ignorance’. We might be deeply ignorant because we are convinced by a false belief, or because we are oblivious to something of moral importance, or because we take some practice so utterly for granted that we do not realise it requires moral interrogation. Where we are deeply ignorant, we might have no notion of the moral stakes or potential consequences of our ignorance, and therefore no ability to guide our inquiries on the basis of these stakes. As I will discuss in what follows, the problem of deep ignorance therefore has serious implications for the demandingness of the obligation to know.

### 2.3 The problem of ‘deep ignorance,’ or the unknown unknown

Donald Rumsfeld famously proclaimed that there are ‘known unknowns; that is to say we know there are some things we do not know. But there are also unknown unknowns—the ones we don’t know we don’t know’. He added: ‘it is the latter category that tend to be the difficult ones’, and that also seems true within this debate. In his argument, Guerrero focuses on known unknowns: to be able to recognise, and therefore abide by, the higher epistemic standard he proposes implies that the wrongdoer is aware not only of their ignorance, but that their acknowledged ignorance pertains to a matter with potential moral consequences, and has some sense of the seriousness of those consequences.

Many of the cases in the literature on culpable ignorance concern examples of known unknowns: i. e. agents who are aware of the moral significance of their decisions; whether these are large or small, obvious or less-obvious. To take one example: Jan Willem Wieland, in an article focused on responsibility for strategic ignorance, focuses on the case of a consumer who must choose between two t-shirts: one is cheaper than the other, but it was produced under unfair and exploitative circumstances (which the strategically ignorant agent prefers not to know).<sup>8</sup> This agent is responding to a known unknown, which is why the option to remain ignorant would be strategic. But note: the fact that there is anything morally significant about buying a t-shirt (that there are potential moral costs to

---

<sup>8</sup> Wieland (2017b, p. 4479). Let us allow the stipulation that buying the cheaper t-shirt under these circumstances is wrong.

others regarding your choice) is precisely the sort of thing that someone could fail to realise. This is one of countless examples of potentially disguised moral stakes: I could fail to realise that I live in (and contribute towards) a deeply unjust society; I could unknowingly continue to use products that are destroying the environment; I could fail to realise that my remarks are harmful or offensive; I could fail to realise that my country is implicated in human rights abuses; I could even fail to realise that I don't have certain entitlements, or that I owe more to others than I recognise. Where we are aware of our ignorance, and the potential repercussions of our ignorance, our obligations to inquire seem relatively straightforward. But what are we to make of these cases of deep ignorance: where an agent doesn't know what they don't know; where they are ignorant even of their ignorance?

One route is to claim that agents are always non-culpably ignorant when they are unaware of their ignorance.<sup>9</sup> On this interpretation, obligations to know only arise when we consciously realise that we have something of moral importance to deliberate about. So the shopper who manages to stay ignorant about the existence of sweat shops, child labour, and other forms of labour exploitation would thereby be blameless. Even morally ignorant agents, who take their morally-ignorant racist or sexist beliefs utterly for granted, might be excused by this approach. (Or indeed, an ancient slaveholder: cf. Rosen 2003). In some respects this suggestion would render blameless the worst agents: ignorant, and even wilfully ignorant agents, who have not even undertaken the most cursory inquiries into the moral nature of their conduct, or who, while wondering if they are implicated in something of moral significance, immediately assure themselves that they are not. In many cases (especially where the relevant information was readily available) such agents seem to be particularly guilty of evincing insufficient regard or even ill will in their ignorance; from the vantage of moral concern, they therefore seem especially blameworthy.

Much of what morally informing ourselves entails is precisely the realisation that there are moral stakes involved in our conduct: to discover, where feasible, the contours of our moral lives; to find out where we are complexly implicated: socially, politically, ethically, and environmentally. This can arguably also

---

<sup>9</sup> The volitionist account of moral responsibility, for which conscious control is a necessary condition for blameworthiness, would make a claim of this sort. Note that while volitionism excludes a great many agents from blame for their wrongs, it is not *inherently* less demanding than alternate views: after all, it would demand a great deal from agents who realised they had significant obligations to know.

apply to less factual and more normative assessments.<sup>10</sup> In a separate paper Guerrero (2017, p. 216) ventures, along these lines, that ‘one must pay some attention to whether there is even a moral issue to consider; one must engage one’s moral faculties and “wish to know” one’s moral duty regarding engaging in some particular conduct or not’. It seems, then, that our obligations to know cannot only come into effect when we have determined that there are moral stakes at play (and their strength), but must also include a general duty to discover such stakes, insofar as we are able, so that we may then deliberate accordingly.<sup>11</sup>

Once acknowledged, this point has serious implications. As we will see in the second half of this paper, many efforts to constrain our obligation to know within our epistemic situation—and evade the emergence of an over-demanding obligation— tacitly rely on the idea that our obligations to know only concern known unknowns, or cases where agents are aware of their ignorance and its potential consequences. But this ignores the fact that we can also be culpable for deep ignorance, especially when the information to remedy this deep ignorance is readily available. This leads us to the relevance of an agent’s epistemic situation in determining what they should know, and which ignorance they are culpable for.

## 2.4 The epistemic situation

So far I have been emphasising the ways in which we can be blameworthy for our ignorance. But of course ignorance is also sometimes a genuine excuse: we are often *not* blameworthy, even for our morally-relevant ignorance, and even when we cause great harm within our ignorance.

As we have seen, there are different ways of characterizing how moral fault for ignorance arises. But a point of relative consensus in debates about ignorance involves the importance of an agent’s epistemic situation in determining

---

**10** I do not share the confidence of those philosophers who hold moral truths to be ultimately self-evident. We often have to engage in moral debates to find our bearings, and we certainly have to do so when it comes to emerging social norms and meanings, or to understanding the ways in which our behaviour might be received by others (especially where our utterances or behaviour might be hurtful or harmful in ways we would not have anticipated in the absence of inquiry).

**11** Vanderheiden (2016, p. 305) also rejects the view that we are blameless for ignorance we are unaware of, writing: ‘No coherent obligation to know can apply only to those already knowledgeable and be defeated by ignorance so profound it cannot recognize its own existence ... lest the appeal of ignorance as an all-purpose antidote to culpability for harming be heightened’.

their culpability. Our notion of blameworthy ignorance cannot be all-purpose; it responds to different times, contexts, opportunities, capacities, and availabilities of information. By and large this feels commonsensical: what we ought to know is related to what we *can* know, and what we can know is related to our epistemic situation. When we hold people responsible for their ignorance (or their failures to act based on ignorance) we routinely appraise how easy or difficult it would have been for them to know what they failed to know.

There have been various ways of elaborating on *why* one's epistemic situation is morally significant. It seems plausible that difficulty can impede the inquiry of a sufficiently concerned, or even a maximally concerned, agent. Where the relevant truth is unknown, particularly esoteric, or skilfully hidden from an agent through deception or manipulation, we might allow that their ignorance is not incompatible with sufficient concern. A different approach concerns our 'reasonable expectations' of what an agent should know, where these reasonable expectations are sensitive to an agent's epistemic situation (see FitzPatrick 2008). Under certain epistemic circumstances, one could remain ignorant despite having inquired extensively, and without having been reckless or negligent in the management of one's morally-relevant beliefs.

Let me introduce a specific case, for reference. This case (taken from Bradford 2017, p. 180) features Samuel, who has failed to recognise that:

The vegetables he is buying were harvested by indentured slaves who are horribly exploited to no good end except the big wigs' profits. Buying these vegetables is wrong if this is true: suppose the big wigs are closely tracking sales and will increase the harshness of the working conditions if they go up.

This example might seem fanciful, but unfortunately in many respects it is not. Some of our seemingly innocuous choices are nefariously implicated: in the displacement of people, the destruction of environments, the spreading of disease, the exploitation of labour. Indeed, a real version of the Samuel case is offered by 1950s potato farming in Bethal, in apartheid South Africa, which relied on prison labour and took place under circumstances so deprived and brutal that they were sometimes fatal; the more demand for potatoes increased, the more incentive there was to charge people in Bethal with apartheid pass law offenses, so as to grow the free labour force. Buying these potatoes, then, was to perpetuate and exacerbate this extreme injustice.

So is Samuel blameworthy for his ignorance, and the wrong of buying the vegetables? Bradford introduces the case of Samuel to consider the extent to which epistemic difficulty mitigates his responsibility for his ignorance. And indeed, on the surface, much seems to turn on how easy it would have been for Samuel to find out this essential moral context, and to recognise his

wrongdoing. An important factor here concerns whether this information was available to him: if there was no way for Samuel to know about these repercussions, his ignorance is clearly innocent. Where the information was available, his ignorance becomes more suspect. But even here, it cannot merely be theoretically available, but must be genuinely accessible to Samuel in practical terms: he would need the opportunity to access it, where this opportunity involves the information being acquirable given both his capacities and resources.<sup>12</sup> If Samuel could not read, for instance, it would change what he had opportunity to know, or if he was living in a totalitarian regime where information was meticulously controlled. Or if Samuel could only have found out via a prohibitively costly membership to a specialised database. Furthermore, if finding out had been dangerous, or involved significant sacrifice, this would also alter our understanding of how accessible the information genuinely was to Samuel, and the extent of his culpability for his ignorance. In the South African case, it would have been difficult to know the circumstances under which Bethal potatoes were produced before 1959, but in this year a widely publicised potato boycott brought attention to the issue: should Samuel have ignorantly bought the potatoes in 1958, we would therefore make a very different calculation about his blameworthiness than if he had ignorantly bought them in 1960.

We see then that there is an essential interplay between the fact of an agent's epistemic situation, and moral evaluations concerning whether they have acted with vice, or insufficient concern, or whether they have met reasonable expectations in their efforts to inform themselves. There is an important sense in which the parameters of our moral responsibility for what we know are therefore established by facts about our epistemic situation: what information was readily available and accessible to us, and required little sacrifice to find out. For the remainder of this paper, I will explore what this relationship implies for the epistemic situation enabled by easy access to the internet, which is the situation many of us find ourselves in. I will argue that our epistemic situation leaves us in an irresolvable dilemma regarding the demandingness of our obligations to know, and that it also leaves us newly unable to morally appraise the ignorance of others (and even, to an extent, our own).

---

<sup>12</sup> Cf. Wieland (2017c, pp. 743, 749). Wieland maintains that one can have opportunity without being aware of it, provided you *could* become aware of it. Another important part of Wieland's investigation is the impact of 'exceptional' information, where 'agents facing exceptional evidence are *blamelessly* unaware ... while lazy agents are *culpably* unaware'. Where information comes from a familiar source, and hasn't been adjusted unexpectedly, it is not exceptional in this sense.

## Section 3

### 3.1 The dilemma posed by our epistemic situation

The proliferation of internet technologies probably constitutes the greatest transformation of our epistemic situation in the history of humanity. This transformation is so variable and complex that it is not the sort of thing we are ever going to have a firm grip on. But that the change is overwhelming is indisputable, and it requires us to think anew about many fundamental questions in moral epistemology.

The quantity and scope of information we now have easy and immediate access to is unprecedented.<sup>13</sup> What's more, engaging with any single facet of this information mostly comes at a trivial cost: we can often access information for free, and it usually takes little time or hassle to do so. Without so much as leaving my bed, I can read the local news from almost any country in the world; pore over a crowd-sourced encyclopaedia with millions of entries; skim some of the 25 million titles on Google Books; look at the 160 million papers on Google Scholar; study reports from the United Nations, Human Rights Watch, Amnesty International, and hundreds of thousands of other organisations and institutions which endeavour to gather empirical research to guide interventions. I can watch tens of thousands of documentaries, I can sign up for free online courses in every imaginable discipline, and I can expose myself to what can only be described as the infinity of opinion, chatter, accusation and analysis that takes place on social media sites, interspersed here and there with important moral and political debates.

As we have seen: the parameters that establish the information we are morally responsible for knowing (in the sense that we could be blameworthy for failing to know it) are established by which morally-relevant information is easily accessible to us, and which information we had opportunity to gather given our resources and capacities. Straightforwardly, then, this realm has never been larger. We ought to know more now than ever, and (according to these criteria) our ignorance on many fronts is more blameworthy than it has ever been.

---

**13** Of course, the internet has not only been an extraordinary source of information, but also an extraordinary source of misinformation. The complications posed by poor quality information and misinformation are important: where information will have little chance of remedying our ignorance (and might even exacerbate it) this of course has implications of our obligations to engage with it. I am largely setting these complications aside for present purposes, and focusing on our obligations towards the abundance of reliable, good quality information within our epistemic situation, which could potentially remedy (rather than exacerbate) our ignorance.

This is not merely academic: the charge of ignorance is endemic, and we routinely see people caught up in scandals of what they ‘should have known.’ We can be, and often are, accused on account of our ignorance and considered blameworthy for it, and the accusation often turns precisely on the easy availability of the relevant information; how simple it would have been for us to know, if we’d only cared to know. Even beyond our practice of holding one another responsible, it might be that we are (in an objective sense) made profoundly more blameworthy by our epistemic situation. ‘They were thus *capable* of living lives of virtue’, writes Dennett (1986, p. 144) of our information-scarce ancestors. ‘Of a virtue that *depended on* unavoidable ignorance’. Perhaps we have genuinely been separated from this virtue, and our lives now manifest vice to a degree unknown before.

A paradox of sorts seems to emerge within our epistemic situation. In many individual cases the expectation that we ought to gather morally-relevant information has never been more reasonable or less demanding: the requisite information is often readily available, and freely accessible. Yet within our context, which is defined above all by *quantity* of information, it is misleading to allow our intuitions regarding these individual cases to guide the parameters of our larger responsibility. Once we consider our full epistemic situation, the demandingness of the obligation to know is crushing. When we dwell on everything that we could ‘easily’ know— about the myriad networks of economic and political exploitation that we’re implicated in, the harms and suffering we’re contributing to or failing to address, the environmental impacts of our choices, about emerging moral, political and social debates, about distant injustice and suffering, about the nuances of our political participation (and so on and on)— we see that the necessity to engage with all of it would quickly be overwhelming. Within our epistemic situation of super-abundant, easily accessible information, the qualifiers concerning the availability of information and the opportunity to gather it, which once forgave us a great deal of our ignorance and radically constrained what we were obligated to know, now excludes less than ever. Under these circumstances the obligation to know is radically over-demanding.

I will soon look at a range of attempts to escape this dilemma, and I will argue that none succeed satisfactorily. In each case, we are left unable to describe an obligation to know that avoids over-demandingness when considered at large, but which is nevertheless sufficiently demanding (and not amorally lenient) in individual cases. We are trapped, then, between over and under-demandingness when we try to describe any particular parameter for the obligation to know within our epistemic context.

Importantly, my focus here is on information that is of relevance to our moral lives; it is in this sense that I have isolated a *moral* obligation to know.

Obviously much of the information that is now available to us does not fall into this category, but our epistemic situation has also profoundly expanded the scope of potentially morally-relevant information we have access to. To bring the moral stakes of this conversation to the fore, let us imagine a series of agents who have committed some wrongdoing in ignorance.<sup>14</sup> Perhaps they made some joke or remark which they ignorantly thought was acceptable, but is in fact harmful or offensive to others in a manner that makes it wrong (transphobic, homophobic, racist or sexist), and which they would have recognised as such if they had only been paying greater attention. Perhaps they are blithely unaware of some massive injustice within their society, or perpetuated by their country abroad, which again they could have realised if they had engaged with the relevant easily-available information.<sup>15</sup> Or perhaps they are continuing to buy a product, the use of which is causing great harm— to people, animals or environments—which again they could have easily determined. (Note that in each of these cases the agents could also have been deeply ignorant).

These are precisely the forms of ignorance that we routinely consider blame-worthy. We respond to ignorant bigoted remarks with blame (and, increasingly, shame), and feel little sympathy for the agents who say they ‘didn’t realise’ or that they ‘didn’t mean’ to offend. We recoil from individuals who claim ignorance regarding the injustices that were perpetuated within or by their societies; indeed, often we interpret this kind of obliviousness as tantamount to sanctioning or condoning the injustice in question. We blame people for the harm they perpetuate in their behaviour and consumption, and often think they ought to know better, or that they have morally failed in their ignorance. It is in this sense that (as I remarked earlier) this question is not merely academic: we are constantly involved in inter-personal evaluations concerning one another’s ignorance, and efforts to determine the moral significance of this ignorance.

With these stakes in mind, let us turn to the various attempts to escape the over-demandingness which emerges from our obligation to know within our present epistemic situation.

---

<sup>14</sup> Let us grant these wrongs for the sake of argument.

<sup>15</sup> At an extreme this can apply to those who allegedly failed to recognise the moral seriousness of Nazism, or of apartheid policy in South Africa. Nomy Arpaly and Timothy Schroeder use the example of Paul Feyerabend in Nazi Austria who ‘unblinkingly admits his lack of moral concern at the time when he says that during “all these events” it never occurred to him to inquire further, for “the idea that the fate of any single human being was in some way connected to my own existence was entirely outside my field of vision”’(Arpaly and Schroeder 2014, p. 243).

### 3.2 Ought implies can

A first response to this dilemma might be to argue that we do not have an obligation to keep abreast of all of this easily-accessible information, because we simply cannot, and ought implies can. Questions of impossibility are separate from questions of demandingness, and staying abreast of everything which might have bearing on our moral conduct would take more time than our lives allow, and would probably overwhelm our processing capacities, especially given how often information of this sort is contested, disputed and revised.

This claim would quickly dispel any obligation to engage with *all* of this information, and where an issue is especially complicated and the information involved is particularly difficult to navigate, we might also be forgiven our failure to arrive at the right belief. But it is nevertheless quite easy to constrain the obligation to the realm of ability, and within this realm the problem of demandingness persists. The ‘robust obligations to know’ that Steve Vanderheiden endorses are sensitive to our ‘cognitive capacity to assimilate accessible facts’ (2016, p. 298). He argues that ‘cognitive limits on the ability to know must allow for some excusable ignorance in the case of persons’ (2016, p. 306). However, insofar as we have *not* exceeded our cognitive capacity to assimilate the accessible facts, this excuse would not apply.

Furthermore, it is clear that the impossibility of the full requirement would in no way excuse us from failing to do anything at all. We could easily do more, and with effort a great deal more, without outstripping our resources of time or our mental capacities to assimilate and process information. Insofar as my efforts in moral inquiry enabled me to avoid unknowingly perpetuating *any particular* wrongdoing, every such discovery would constitute a moral contribution towards the good. The obligation to know can therefore continue to overburden us even within the framework of our genuine abilities.

### 3.3 Prima facie obligations

It might then be pointed out that the obligation to know is not proclaiming to be all-things-considered; rather, it is a prima facie obligation that can, and often will, give way to more pressing obligations. We might think a doctor has obligations to keep up to date with the medical literature, but if she finds herself in the midst of a terrible war or epidemic, and is called upon to put all of her energies towards emergency medical treatments, it seems quite clear that her medical journals may pile up unread, and her conferences may pass by unattended. All of us have obligations which plainly trump our more general obligations to know. This is

acknowledged even by those philosophers who take a strong line on our blame-worthiness for failures to inquire. Wieland, for instance, writes that ‘concerned agents [need] not spend energy on preventing X if they have more important things on their mind (that is, things which are *in fact* more important, not things that the agent *thinks* are more important’ (2017b, note 10, his emphasis).

But how easily should other obligations triumph over our obligations to know? It could be argued that such obligations should not be easily overridden (which seems to be what Wieland is implying above). In this case, we are only excused from our obligation to know when it competes with a moral obligation which is plainly more important. But it seems that a great many of our interests and concerns, and what we take to be significant and devote our mental and other energy to, cannot possibly compete for warranted priority. Very little that occupies my time can claim to be *in fact* more important, morally speaking, than determining the complicities and repercussions of my actions, or concerning myself with suffering and injustice in the world. It seems, then, that despite acknowledging its prima facie nature, our epistemic situation still leaves us with an over-demanding obligation to know, since the only relief from the obligation would be tending to another more serious or urgent obligation.

Alternatively, it could be argued that other interests and obligations should quite easily override our moral obligations to know. Bradford takes an approach along these lines: she holds that blame is mitigated when knowing comes at the expense of other goods, and she is generous regarding what constitutes these goods, including ‘the good of living as the author of one’s life’ (2017, p. 189). This seems to suggest that we can have an obligation to know which is not over-demanding since ‘the good of exercising one’s own agency... outweighs the bad of the wrong’ (2017, p. 196). But this is only plausible where the bad of the wrong is trifling or moderate. One could hardly say, where some ignorant agent had perpetuated a significant wrongdoing (where Bob does not investigate the toxic waste, for instance), that they are blameless for their ignorance on the basis that they were living as the author of their own life. Such a good ought often to give way, including giving way to the obligation to know, and indeed it would have to if we are not to establish an obligation to know that is so feeble it would even permit harmful negligence. It seems, then, that when we allow other interests to prevail too easily, the obligation to know fails to be demanding enough, but when we set a high bar on which obligations ought to prevail, the obligation to know is over-demanding.

To navigate this tension, Bradford suggests (following Guerrero) that the necessary effort in inquiry must correspond to the moral stakes involved. Where the moral stakes are very low, we can prioritise our own interests, but where they are high, we are obliged to make the effort to inform ourselves. But the problem with

this position arises once more, since in many cases *it is precisely the stakes that are unclear*. Let us return to Samuel: he doesn't know how serious the harm he's contributing to is, so he is not able to appraise whether his other interests and obligations—and the good of living as the author of his own life—should outweigh his obligation to inform himself. Much of the strength of this obligation will depend on what is then discovered. How can an agent know whether, in a particular case, other goods and obligations should prevail, unless they are already cognisant of the moral stakes involved? Given these intractable uncertainties, the *prima facie* nature of the obligation does not seem to spare us from the dilemma of demandingness.

### 3.4 Common knowledge

A different approach to the demandingness problem is to suggest that our actual obligation is only to keep abreast of 'common knowledge.' We are spared any obligations to the infinity of 'available' morally-relevant information, and need only stay alert to information that it prominently broadcast, or has seized national or international attention (what Dennett (1986, p. 145) describes as 'a small, shifting core of what might be called "temporarily famous" common knowledge'). So on this interpretation, there would be no excuse for someone using plastic straws, now that there has been a widely publicised straw boycott, though perhaps one would be forgiven for not realising how environmentally costly it is to produce almond milk.

This seems to be a fairly good description of how we behave in practice. And when we think about the overwhelming nature of our epistemic situation, we are most likely thinking of the endless barrage of breaking news, or headlines, which we feel obliged to attend to. We often conflate our obligation to know with our obligation to stay up to date with this stream of common knowledge, which nowadays feels more like a deluge.<sup>16</sup> But it is important to note that while common knowledge is particularly readily accessible (we usually encounter some of it socially, even if we make no effort to keep abreast of it for its own sake), it is by no means the only information that we have ready access to, even if we might gravitate towards it.

When we bear this in mind, it is clear that we cannot constrain our obligations to know merely to common knowledge. Indeed, this approach seems

---

<sup>16</sup> The question of whether we can disengage with the news, or with 'common knowledge,' is separate from the more general question of our obligations to know. One could argue that we can disengage from much of the common knowledge—and even that we ought to—while retaining robust obligations to know in other spheres.

morally bankrupt. Crucially: it mistakes what we *are* paying attention to, for what we *ought* to be paying attention to; and two categories come apart, often enormously.

When we dwell on the nature of the stories that routinely command our collective attention, it is patent how little relationship there is between them, and what ought to be guiding our moral inquiry, and our obligations to know. In order to compete for attention, news often responds to the sensational or entertainment value of information, rather than its moral value. We are consumed by shame campaigns, scandals, public disgrace, celebrity meltdowns, and political soap operas. It has also always been the case, and remains so, that what demands much of our attention in ‘current affairs’ is more likely to be the exceptional and unusual, rather than the slow-burning suffering which is constant and therefore never ‘new.’

It is clear that common knowledge often neglects things we ought to know about, that have bearing on our moral lives, and so we should not cease our inquiry on a few front pages and evening broadcasts, but ought to look in less obvious places, and consider the quiet tragedies along with the loud ones. It seems, then, that any attempt to rein in our obligation to know by restricting it to common knowledge would be fundamentally misguided.

### 3.5 Individual versus collective obligations

Another route is to distinguish between individual and collective obligations to know. We might argue that states (for instance) have a greater obligation to know than individuals do, and therefore that the state lifts some of the burden of the obligation from the individual. Vanderheiden takes an approach of this sort. As we saw above (in the section on ‘ought implies can’) Vanderheiden argues for robust obligations to know in the case of individuals, but he is simultaneously cognisant of our limited capacity to assimilate information. He suggests that states do not have a similar excuse: ‘states and other large-scale organisations have a much greater ability to process information than do individual persons, and are as a result more circumscribed in their claims to excusable ignorance’ (2016, p. 306).

Vanderheiden is not explicitly concerned with questions of demandingness, and he draws the boundaries of our individual obligations close to the limits of capacity. But we might separately imagine a stronger version of this position, which argues that the state ought to relieve individuals of a great many of their obligations to know, and thus generates a more modest individual obligation which is not over-demanding.

If one lived in a state that was fulfilling these obligations, and pursuing moral ends for their own sake, perhaps one would only need to make sure the state was taking these various obligations seriously, and insofar as they were, you would not have to concern yourself a great deal further. Unfortunately for most of us though, this sounds utopian. And it is hard to know what one's individual obligations become in circumstances in which the state is renegeing on its obligations to know and to act. Insofar as we live in states which do not pursue moral ends for their own sake, we might be obliged to campaign for the state's resources and attention to be directed at certain moral issues. What's more, often it seems that rather than being relieved of our obligations to know by the state, we actually have more obligations to know which *concern* our states, and the moral problems they are contributing to (rather than resolving) in our names. Appealing to the epistemic obligations of states therefore does not seem to resolve the threat of demandingness when it comes to our individual obligations to know.

### 3.6 Aggregation

Finally, we might consider an aggregative rather than an iterative approach to our obligations to know, as a way of reconciling these obligations with the overabundance of available information in our epistemic situation. The over-demandingness of our obligation to know, as I have been describing it, is generated by a process of iteration: in every case of harmful ignorance, we see that the relevant information was easily available, and therefore that the agent's ignorance is culpable. But if we opted for aggregation instead of iteration, we would not look at every individual case in isolation, but rather at the broader picture of what someone has endeavoured to know. In this case we might say that we have fulfilled our obligations where we know *enough*, where this is established by some aggregative threshold. Where we have met this threshold, we are forgiven our ignorance on other matters, and excused from the wrongs we commit within this ignorance.

Initially this suggestion seems promising: it provides a way to maintain our moral obligations to know, but doesn't topple us into a kind of boundless or all-consuming obligation, which seems inherently over-demanding. What's more, this approach is comparable to the way we sometimes think of our moral obligations in other overwhelming spheres, such as charitable donations; when we donate money, we cannot possibly donate to *all* worthy causes, but only to some.<sup>17</sup>

---

<sup>17</sup> I thank an anonymous reviewer for suggesting this comparison in favour of aggregation.

But aggregation generates problems of its own. In a case where an agent committed wrongdoing X in ignorance (despite a wealth of available information that they could easily have consulted), there is something uncomfortable about pointing to everything that they *do* in fact know as mitigating. So if we take Samuel again, and imagine that the information regarding the vegetables is widely and readily available, can we really say (in not holding Samuel accountable for his ignorance and its repercussions) that he is excused because he separately knows a great deal about the endangered Leopard Toad, and the scourge of micro-plastics in our environments? I think many of us would be tempted to hold Samuel responsible, regardless of what else he happens to know.

While an aggregative approach might make sense when considering information we know for its own sake, it does not seem like an adequate way to appraise the information we know *in order to inform our moral conduct*. When it comes to wrongdoings and harms I ignorantly perpetuate, my culpability for my ignorance does not seem alleviated by counterfactuals concerning where I am properly informed, especially where the information to remedy my harmful ignorance was readily available to me.

To develop this idea, let me return to the agents I introduced earlier: who ignorantly made bigoted remarks, or overlooked prevalent injustice, or perpetuated great harms (despite access to information which would alleviate this ignorance). If we attempted an aggregative route, the threshold of how much we ought to know would presumably have to be set quite high, so as not to be amorally lenient. Let us grant, then, that these individuals do know a great deal about other moral issues; enough to qualify on an aggregative interpretation of our obligation to know. Depending on the ignorance in question, we might *still* feel that their ignorance is simply not exculpated by facts about their other morally-relevant knowledge.<sup>18</sup> Is there any amount that someone can know about other moral issues which would forgive them for being a climate change denier, or for ignorantly making racist or sexist comments, or for pouring chemical X down the drain and contaminating a local lake?<sup>19</sup> Or even for failing to compute the moral heinousness of Nazism? And where we feel that these things ought to be known, irrespective of what

---

<sup>18</sup> In this respect the analogy to charitable donation is misleading. Ordinarily, we treat charitable donation as supererogatory, but this is not the case when it comes to our ignorant wrongdoings. In appraising the culpability of these ignorant agents, it does not seem appropriate to say: ‘well, so long as you know enough other things’ (in the way it might feel appropriate to say: ‘well, so long as you are charitable in other spheres’).

<sup>19</sup> The final example is taken from Vanderheiden (2016, p. 303).

else these agents know, we seem to give up on the spirit of aggregation: a consequence of aggregation is precisely that serious ignorance about serious wrongs will be rendered blameless despite the easy accessibility of the information necessary to rectify this ignorance.

Separately, it is unclear to what extent aggregation will genuinely limit our obligations to know. To return to the comparison with charity: if we asked whether someone was donating money to help in urgent famine relief, and they responded that they were not, but they had already given a substantial donation in support of the Westminster Dog Show, we might well feel that their support of the Westminster Dog Show did not excuse their failure to help in famine relief, given that the one cause is so much less worthy than the other. Similarly, we might think that less worthy moral knowledge does not forgive ignorance of more important moral knowledge. We might be tempted to appraise the worth of what people have dedicated their moral-epistemic energy to, and weigh it up against the worth of what they could have dedicated it to; perhaps we think Samuel ought to have known about the racist and exploitative conditions in which his produce was harvested since it is so much more important, morally speaking, than the Leopard Toad. But if, in order for aggregation to be successful, we are required to appraise the worth of the various moral knowledge we are pursuing, and weigh it up against that which we are neglecting, then we hardly seem to have escaped the enormity of our burden: an appraisal of this nature, within an over-abundant epistemic situation such as our own, would be very demanding.<sup>20</sup>

Perhaps we could instead suggest that you need only focus your epistemic energy on a few of the big moral issues (e. g. climate justice, poverty, inequality, animal welfare, discrimination). This might spare us the worry that we could squander our epistemic energy on relatively trivial moral matters, at the expense of far more urgent issues. But if we were to take this approach, would we then be permitted to remain largely ignorant on the others, and the ways in which we might be implicated in them? Again, I think this is an uncomfortable suggestion, and certainly out of keeping with our usual modes of assessment. Where someone is so preoccupied by climate justice and animal welfare, for instance, that they take little interest in matters of poverty, inequality, discrimination and injustice, we usually treat their ignorance on these other matters as blameworthy, or at least as reflecting poorly on them. And where someone is not only ignorant, but also ignorantly contributing to wrongs or harms, then this seems all the more obvious.

---

<sup>20</sup> I thank the editors of this issue for this suggestion.

It is separately important to note, in terms of blameworthiness, that this aggregative approach would recommend a profound overhaul in how we appraise people for their ignorance. It is certainly not common practice for us to respond to an agent's harmful and rectifiable ignorance by undergoing a thorough assessment of their other morally-relevant knowledge and concerns. What's more, it is quite unclear how we would go about making such an appraisal. Where would we begin, in trying to determine if Samuel was innocently ignorant of the repercussions of his purchases, in our assessment of everything else of moral relevance he's informed himself about? It is in this sense that our present epistemic situation leaves us uniquely unable to morally evaluate one another's ignorance. Where we are excused by the scarcity of information, it is quite clear which ignorance we are not culpable for: where morally-relevant information is easily accessible, we ought to know it, and where it is not, we are more easily forgiven. But where we are excused by the *abundance* of information, the situation is profoundly different: there is no way to tell what, precisely, we ought to have known amongst all that we could so easily have found out. We might find, on these aggregative terms, that much of the ignorance we would currently take to be blameworthy should instead be blameless.

## 4 Conclusion

It seems, then, that all of these attempts to evade the over-demanding obligation presented by our epistemic situation are ineffective or troubled: where we manage to describe a less burdensome broader obligation, we land up with over-lenient assessments with regards to individual cases, which don't offer sufficiently robust obligations to know, and which don't seem to defend the good of evading avoidable harm, or recognising where we are complicit or otherwise implicated.

But the dilemma of over-demandingness within our epistemic situation persists, and it is perhaps precisely on account of this genuine dilemma that we often feel a deep moral inadequacy in the midst of all that there is to know. Vanderheiden writes that 'excusable ignorance ... can result from too much as well as too little information being available' (2016, p. 304). As we have seen, something about this must be right: not only because of our limited capacity to process information, but also because too much information can present us with an over-demanding moral obligation to know, even within the bounds of our capacities. There is a limit to what we can discover and what we can process

adequately, and there is also a limit to what we can engage with decently and in a manner that is expressive of moral concern: at its widest circumference, this limit emerges on the basis of our capacities, but it also emerges, at a less burdensome reach, on the basis of demandingness.

Unlike the excuse of too little information, however, we are left in a position where it is very difficult (if not impossible) to appraise excusable ignorance in the face of *too much* information. This might even apply, to some extent, to assessments of our own harmful ignorance: some of which must be excusable, but much of which is not. How shall we grant where the excuse should apply to ourselves, or where we should feel guilty or remorseful? How should we determine when we know enough? With regards to our obligations to know within our present epistemic situation, we are left having to manage these vague boundaries for ourselves: to allow that we needn't demand too much from ourselves, but to acknowledge that we can also demand too little.

**Acknowledgements:** Many thanks to the organisers and participants at the 'Demandingness in Practice' workshop, who provided invaluable feedback on an early draft of this paper. Thank you also to the editors and reviewers of this issue, and particularly to one anonymous reviewer who went into such generous detail in their comments; my debts to you were too many to acknowledge within the paper.

## References

- Arpaly, N. and Schroeder, T. (1999). 'Praise, Blame and the Whole Self', *Philosophical Studies* 93 (2): 161–188.
- Arpaly, N. and Schroeder, T. (2014). *In Praise of Desire* (Oxford: Oxford University Press).
- Bradford, G. (2017). 'Hard to Know', in P. Robichaud and J.W. Wieland (eds.). *Responsibility: The Epistemic Condition* (Oxford: Oxford University Press), pp. 180–198.
- Dennett, D. (1986). 'Information, Technology, and the Virtues of Ignorance', *Daedalus* 115 (3): 135–153.
- FitzPatrick, W. (2008). '[Moral Responsibility and Normative Ignorance: Answering a New Skeptical Challenge](#)', *Ethics* 118 (4): 589–613.
- Guerrero, A. (2007). 'Don't Know, Don't Kill: Moral Ignorance, Culpability and Caution', *Philosophical Studies* 136 (1): 59–97.
- Guerrero, A. (2017). 'Intellectual Difficulty and Moral Responsibility', in P. Robichaud and J.W. Wieland (eds.). *Responsibility: The Epistemic Condition* (Oxford: Oxford University Press), pp. 199–218.

- Robichaud, P. and Wieland, J.W. (eds.). (2017). *Responsibility: The Epistemic Condition* (Oxford: Oxford University Press).
- Rosen, G. (2003). 'Culpability and Ignorance', *Proceedings of the Aristotelian Society* 103 (1): 61–84.
- Smith, H. (2014). '[The Subjective Moral Duty to Inform Oneself before Acting](#)', *Ethics* 125 (1): 11–38.
- Vanderheiden, S. (2016). 'The Obligation to Know: Information and the Burdens of Citizenship', *Ethical Theory and Moral Practice* 19 (2): 297–311.
- Wieland, J.W. (2017a). 'Introduction', in P. Robichaud and J.W. Wieland (eds.). *Responsibility: The Epistemic Condition* (Oxford: Oxford University Press), pp. 1–28.
- Wieland, J.W. (2017b). '[Responsibility for Strategic Ignorance](#)', *Synthese* 194 (11): 4477–4497.
- Wieland, J.W. (2017c). 'Evidence One Does Not Possess', *Ergo* 4 (26): 739–757.