
IGNORANCE AND THE HARM OF DEATH

Thomas Hofweber
UNC Chapel Hill
hofweber@unc.edu

Draft to be presented at the Aristotelian Society

Comments are much appreciated

May 3, 2026

ABSTRACT

This paper argues that a key part of the harm of death is connected to ignorance. I propose that ignorance in special cases is tied to harm and suffering, and that the ignorance connected to death normally brings with it this epistemic harm. Furthermore, I suggest that this epistemic harm is a significant, although not the only, aspect of the harm of death. This more intellectualist picture of the harm of death is contrasted with a more hedonist picture which focuses on missed pleasures. I conclude with some reflections on what might explain this connection between death, ignorance, and harm, and point to some ideas motivated by at least the metaphor of a narrative conception of the self.

Contents

1	Introduction	1
2	Ignorance and suffering	2
3	Ignorance and harm	5
4	Ignorance and death	6
5	Ignorance and the self	8
6	Ignorance and time	10
7	Conclusion	11
	Bibliography	11

1 Introduction

Whether death is a harm for the person who dies and what precisely makes death bad is puzzling. Although much has been written to resolve this puzzle, I can't say that I think the issue has been resolved. This paper also won't solve the puzzle, but it hopes to make a contribution to move the debate into the right direction and hopefully also to get closer to a proper resolution to the puzzle why it is bad to die for those who die. Although the standard Epicurean argument that death is no harm at all for us, since we are not going to be around to suffer that harm, has been successfully refuted, in my opinion, by deprivation accounts of harm, nonetheless, these deprivation accounts do not properly address what the harm of death really consists in. The Epicurean is mistakenly too internalistic about harm, essentially insisting that any event that is supposed to harm us must be appreciated or experienced by us to in fact harm us. The externalist alternative can accept that events at different times and far away can harm us without our ever experiencing them, even events at times when we are not alive. And a deprivation account furthermore adds that things that would have happened to us can contribute to our harm or benefit, extending the externalism beyond the factual to the counterfactual. And although this seems correct as far as it goes, it also seems a bit vacuous, especially when it comes to the issue of the harm of death. In a sense, all harm can be understood in such a way. For example, when I twist my ankle, then the harm associated with this event could be seen as a weighted sum of all the good and bad things that did happen because of it as well as the weighted sum of the good and bad things that would have happened had I not twisted the ankle. There is the pain of the injury, the fact that I made a friend at rehab, the fact that I missed out on going to the concert, and so on, all of which need to be added up and properly weighed. The harm of death is just like that, except that there is only a more limited sum that corresponds to what that did happen, since nothing experiential gets added in any more. Instead, the values corresponding to what didn't happen because of my death but would have happened otherwise, play the major role.

All this is just a special case of harm in general, a special case where a sum has fewer components, but with nothing uniquely tied to death. And because of this the account seems to be lacking, since the case of death plausibly is special. It is not that the deprivation account is wrong, it just doesn't get to the heart of the problem. It is in a sense too general and high level, leaving the real issue aside, the issue about what it is about death in particular that is bad.

The hedonistic version of the deprivation account of the harm of death, according to which the harm is based on the sum of the pleasures and pains one would have had, seems to me to be especially inadequate to get at the real issue. If I died now I would miss out on the delicious pizza I would have had next week, but to think that the harm of death essentially comes down to a sum of missed pizzas and other missed experiences seems completely wrongheaded to me. I had plenty of pizzas in my life, and although I surely would like to have another one in the future, the harm of death does not seem to me to consist in missed pizzas and other missed pleasures. There is something else there, something tied to the life one has lived coming to an end too soon that must play a larger role in an account, not just that one's pleasurable experiences disappear. In this paper I hope to make progress on what that might be and how it can help us explain why death is a significant harm and what makes it bad for those who die.

I made my dislike of the attempt of trying to solve the problem by combining hedonism and the deprivation account clear, but the issue, again, is not the deprivation account as such. The deprivation account is in a sense correct but not complete. One does not have to restrict oneself to a hedonistic version, focusing on missed pleasures, but one can go broader, considering all other things of value that one would miss. The harm of death is then the weighted sum of all the harms and benefits one would have had if one had not died. Harms and benefits here can be quite general, including pleasures and other things that are not directly experiential: wisdom, friendship, and so on. That general account of harm might well be true, but it also doesn't really help with understanding the harm of death in particular. It at best makes clear that it makes sense to talk about death being a harm for the person who dies, but it doesn't illuminate what that harm is and whether it is a significant harm. We must do better in this regard.

All this also highlights that there is not a unique harm of death, and thus that it is slightly inaccurate to speak of "the" harm of death, as if there was just a single thing that makes it bad to die. Missing out on pleasure like the pizza next week is one harm of death, we can agree, but there are also others that are very different. This suggests a pluralist account of what the harm of death consists in, and it suggests a distinction between the main, principal sources of this harm and the more minor, secondary sources. My own sense is that missed pleasures at least of the kind one has experienced many times before are a source of harm, but not a principal one. This then raises the question what the principal sources of the harm of death are and why it is those.

In this paper I will outline a proposal to make progress on this issue. I will argue that a principal harm of death is tied to ignorance. If we die we won't come to know the answer to certain important questions, questions that we would likely have known the answer to simply by living longer. To defend this position I will first elaborate on a somewhat surprising and neglected connection between ignorance and suffering, then develop the view that a particular kind of ignorance is a key aspect of the harm of death, and finally try to explain why the harm of a certain kind of ignorance is especially connected to death by investigating the relationship between knowledge and the self.

2 Ignorance and suffering

Ignorance contrasts with knowledge. I know who won the last soccer World Cup, but am ignorant about who will win the next one. In general, knowledge is better than ignorance, all things being equal. There might be some cases where it is better to be ignorant, say regarding what one's friends say behind one's back or what secrets family members have. But in general, it is better to know than not to know. It is better to know about science, history, literature, ordinary facts, etc., than not to know about these things. Ignorance is a negative state, while knowledge is a positive one. So far most would agree.

But what is more, ignorance can be suffering for the person who is ignorant, and not just a little bit of suffering, but profound suffering. To say that ignorance is a negative state is simply to evaluate it from a distance. That distanced evaluation by itself does not explain the effect that ignorance has on those that are ignorant for particular cases. And this connection of ignorance to suffering and harm is both puzzling and central for our main topic of the harm of death.

To illustrate the point, we can consider a series of examples where ignorance is connected to suffering to an increasing degree. First, consider a case of inconsequential ignorance. Say I can't remember the name of someone who I briefly met and who I got reminded of just now. I might never meet this person again and they might never come up again, but still, I am bothered by not being able to recall their name. This ignorance leads to a little bit of suffering. It is inconsequential and small, and likely just based on the fleeting desire to know and remember. But the degree of suffering can be much greater in other cases.

To take it one step further, consider a mathematician who struggles to find the answer to a particular question and is tormented by them not knowing the answer. They are struggling with a mathematical problem, they put in their best effort to solve it, and yet, they remain ignorant of what the answer is. This ignorance can make them miserable, even if nothing else hangs on them solving this problem, like their career or employment. Nonetheless, they are engaged in the project of finding the answer and this project being frustrated by ignorance leads to more serious suffering. Here it is not just a fleeting desire to know that is frustrated, but a more substantial project.

But the most compelling cases to illustrate the connection are those tied to interpersonal relationships that are of great significance and ignorance tied to them, especially the case of a missing person. When someone goes missing with no clear indication of what happened to them, where they are, or why they are gone, then this unsurprisingly can bring great suffering to the family and friends of that person. Some aspects of this suffering are not tied to ignorance and knowledge. They are simply the result of friends missing out on good times, being deprived of the company of the missing person, and so on. But other aspects of the suffering are what we can call *epistemic suffering*: they are tied to one's epistemic states, in this case based on the ignorance of what happened to the missing person. Where are they now? Why did they disappear? Are they OK? Simply not knowing what happened leads to suffering. But on reflection, this is puzzling. It is understandable that knowing that something bad happened to them leads to suffering. After all, that means something bad happened to someone close, and that bad thing supports and explains our suffering via our finding out about it. My knowledge of this bad thing might cause my suffering, but it is what I know about that justifies it. My suffering here is not based on my epistemic state directly, but on the bad thing that happened. But why then does the mere ignorance about what happened lead to suffering?

Maybe the explanation goes along very general lines: if one doesn't get what one wants, then one suffers. Being ignorant while wanting to know is a special case of this, and thus ignorance leads to suffering. Although there certainly is some kernel of truth in here, this general connection does not seem to explain the extent of the suffering. Simply because one really wants to know what happened doesn't seem to properly get to the heart of the issue. Knowing what happened to the missing person is not simply something one happens to really want, just like someone might happen to really want that new iPhone. Or to look at it from a slightly different angle: why does one want to know so badly, as opposed to simply want for the person to be fine? Concern about knowledge seems to be about oneself, about one's own epistemic state, rather than concern about the missing person directly. Similarly, ignorance is a state of oneself, not directly about the missing person. Why focus inward, rather than on the missing person?

Let's consider a case in more detail where one's own ignorance substantially contributes to one's suffering. Naturally such a case has to involve ignorance of something rather significant, and a good example of this is the case of a parent whose child is missing. They do not know where the child is, if they are OK, if they are still alive, if they have suffered since their disappearance, and so on. Clearly this ignorance leads to substantial suffering. And also clearly, this suffering is tied to not getting what one wants: one wants to know but doesn't. But the suffering does not seem to be simply based on the fact that one of one's desires isn't satisfied. If one had the option of getting rid of the desire to know, through some medical intervention or drug, then the true source of one's suffering would not be removed, one would just be appeased to it. I suffer because I don't know what happened to my child, not because I am dealing with another case where I don't get what I want. The suffering is thus not merely based on not getting what one wants, even though obviously one does not get what one wants. But then, why does one suffer? This suffering of a parent tied to the ignorance about what happened to their child is puzzling, since can seem to be too closely tied to one's epistemic state and not closely enough to the child.

It might be tempting to think that one's aiming for knowledge is tied to the child after all, since this knowledge might help one intervene and find the child, and thus diminish the harm to the child. And in general there certainly is such a connection. Knowledge supports action and ignorance hinders it, and maybe this action to lessen the harm of the child is what is ultimately the goal, and one's inability to act on behalf of the child is the basis of one's suffering. But on reflection this line strikes me as mistaken. Even in cases where one knows that one can't intervene one still suffers from ignorance. If the child is on a mission to Mars and all communications of the spaceship go silent, then I know whatever happened, I can't do anything about it. Or consider a case where I am prevented from helping do some constraint on me, I am imprisoned and isolated, say, and I know I can't help no matter what, still, I suffer from the ignorance of what happened to the child. The suffering seems to be based on the ignorance itself, not something else that is tied to the ignorance.

To illustrate this further, it is informative to see what happens when one learns that the child is dead and their body has been found. As has been reported in actual occurrences of such a case, this can lead to relief, and it ends a certain kind of suffering, although it clearly does not end all suffering of the parent.¹ But on reflection this should be puzzling. Wouldn't it be better to remain ignorant and to have a small chance that one's child is still alive? Even if they have been missing for a while, there is always a small chance that they simply ran off, hiding out somewhere, maybe living with a different family, or at least that they can be saved later on. That chance disappears once their body is found. What also disappears is ignorance about whether they are still alive and likely also some degree of ignorance about what happened to them. But why does the small chance of the person being alive not outweigh the knowledge we gain about what happened to them? It would seem that no matter how small the chance, it should be better to have that small chance that they are alive and well somewhere than to know that they are dead. Unless, of course, the badness of the ignorance itself needs to be factored in rather substantially. The way to make sense of all this is to acknowledge that ignorance of what happened leads to suffering, and the end of this suffering outweighs the disappearing of a small chance of survival

¹See (Boss, 1999).

for the missing person, at least for the person who suffers from the ignorance. For this relief after finding out what happened to make sense, the positive feature of this finding out must be the knowledge of what happened and the end of the corresponding ignorance, and with it the end of the suffering tied to ignorance.

But why that would be so is still puzzling. Ignorance is just an epistemic state of a person, it by itself does not hurt or feel like anything at all. Why would this ignorance lead to suffering of this magnitude? This asks for some explanation. I mentioned one unsatisfactory option to explain this just above: one wants to know and by finding out one gets what one wants. On the flip side, if one remains ignorant, then one doesn't get what one wants, and not getting what one wants leads to suffering. But again, this does not quite seem to get to the heart of it. If I only get the new iPhone if the missing child is dead, then it would be monstrous to find relief in that they are dead so I can finally get what I really want: a new phone. This would be a misplaced relief indeed, but I don't have the same reaction towards the parent who feels relieve simply from knowing about what happened to their child, and thus getting what they want — to know — even it is not the outcome they were hoping for. Thus the puzzle remains why there is a positive outcome, a kind of relief, from coming to know but forfeiting the chance of the child being fine, no matter how small. Why would an epistemic state of myself outweigh this small chance of an otherwise great outcome? One way could be the enormity of the suffering tied to ignorance in this particular case. But why would there be such significant suffering tied to this ignorance?

Besides focusing on relief tied to finding out, we can also rely directly on reports of those connected to a missing person, and how they are tormented by not knowing what happened. It is ignorance specifically that contributes to suffering in these cases. It is not simply the possibility that something bad has happened, or the chance that something bad has happened, but ignorance about what happened.

To be clear, not all ignorance is suffering. I don't know how many leaves the gingko tree outside my office has, but that ignorance does not cause any suffering in me. We should distinguish *mere ignorance* from *inquisitive ignorance*. The later concerns ignorance about an issue that one is invested in, possible actively trying to investigate and find the answer to. Inquisitive ignorance comes in degrees, not just with regard to how little one knows in the relevant area, but also with regards to how deeply one is involved in the project of finding the answer or how much one is concerned about what the answer is. We can thus say that one has *deep concern* about a question one is very deeply involved in, and about the answer one is presently inquisitively ignorant. It is this deeply concerned inquisitive ignorance that is tied to the harm of death, in particular for special cases, or so I want to propose.

3 Ignorance and harm

Above I connected ignorance to suffering, and suffering is in turn connected to harm. How this last connection goes more precisely is not completely obvious. It is tempting to hold that harm is a broader term which covers suffering as a particular kind of harm. Suffering seems to require an experiential component, something tied to how things feel to someone, whereas harm does not seem to require this. As acknowledged above, on a plausible broadly externalist

conception someone can be harmed by the actions of others even if they never find out about it, and thus they do not suffer from the actions of those actors. Whether this broader conception of harm is correct is not completely clear. It could be, for example, that due to some error or some false belief one falsely takes there to be a harm and suffers as a consequence. In this case one might hold that one is suffering even if one is not harmed. But then again, maybe the ill-placed suffering itself is a harm, and thus suffering is always associated with a harm, even if it is not based on a harm. For us these subtleties won't matter, since we won't be dealing with cases of errors like this. We can thus take harm the more general, more externalist term, with suffering being a subjective experience based on a harm. In epistemic suffering, this suffering is then based on the *epistemic harm* of ignorance, but the suffering adds a subjective experience on top of it.

The notion of epistemic harm is also used in a broader sense than we will use it here. In the literature on epistemic injustice² the notion of epistemic harm is used to indicate any harm to a person tied to epistemic issues. For example, undermining someone's credibility results in harming them, since they cannot communicate what they know as well as before. For us here epistemic harm is more narrowly focused on the harm a person suffers due to ignorance.

The epistemic harm tied to ignorance due to death can be nicely associated with a general deprivation account of harm. I am deprived of the knowledge I would have had had I lived longer, and thus I am deprived of the benefit that comes with this knowledge. The question remains, of course, how great such a harm is and why it is there in the first place. To connect the harm of death to epistemic harm we need to say more. To see that connection, it might help to first think a bit about ignorance about one's own life.

4 Ignorance and death

Normally, we know how our own life is going day to day, simply by living it. But that does not mean we know how our life is going overall, as a unified single thing. How valuable is it to know not just about one's day to day life, but one's life overall? What price should one pay to come to know what life overall one has lived? There is an illuminating thought experiment that concerns just this question: would it be preferable to die in one's sleep or die awake, fully knowing what is happening, suffering through it? In a brief essay (Velleman, 2015), David Velleman proposes that dying awake is to be preferred even if one thereby experiences the pain and suffering of dying. As he says " ... closing the circle of one's life is necessary to seeing it as the particular life one has lived. And I want to know the particular life I've lived before I stop living it — which will entail fully living it up to the very end." (Velleman, 2015, p. 198). This strikes me as very plausible. One wants to see one's life as a whole at some point, and the only point where this is truly possible is at the end. Dying by surprise in one's sleep deprives one of this perspective on one's own life in its totality. It would be easier to go this way, but at a great cost. Now, Velleman does not want to directly connect this thought to the harm of death, as he says in the same essay: "I have lost my grip on the philosophical question about the harm of death." (Velleman, 2015, p. 196). Velleman's reasons for this are complex and I am not sure I agree with him here, but there is a

²See (Fricker, 2007).

clear connection to the harm of death in this thought experiment of dying in one's sleep which I find illuminating and plausible. Let me thus present the basic idea and how I see how it connects to our main topic.

Knowing how one's life ends is an epistemic state, and so is seeing one's life as a whole, from beginning to end, and thus coming to know overall what life one lived. This is a case of knowledge of deep concern, and dying in one's sleep deprives one of this knowledge. A hedonist might think it is worth it overall, since it also deprives you of the unpleasant experiences that come with dying. And although I am no fan of unpleasant experiences, I am with Velleman here. Dying without ever seeing your life as a whole, without ever coming to know what life one lived overall, leaves one's life incomplete. This is made clear by the fact that one does not simply come to know about events that just happen, but rather one is actively involved in making them happen. Living a life involves making this life happen through one's agency, which in turn is controlled in part by one's knowledge and goals. These goals do not just locally guide us to go left or right, but globally to live this way or that way. And that involves finding out whether one lives as one hoped and aspired to, which is a question about one's life as a whole.

These considerations so far support that dying in ignorance can be harmful, in the sense that dying without knowing how one's life was as a whole can be a harm greater than dying while suffering physical pain. But that does not yet support that there is a special connection between dying and ignorance, since dying with one's eye's open, so to speak, avoids the harm of ignorance about one's own life's completion. That ignorance can be avoided at a cost, but the issue this points to is broader. We are not merely harmed by ignorance about our own life in death, but ignorance about many other questions that are not directly about ourselves.

During our life we are generally ignorant about the answers to many important questions about people we are close to, our community, and about humanity in general. These questions are of especially deep concern, although certainly not the only ones of significance. But finding out what happens to us, to our family, to humanity, how our lives individually will go and how our community and species in general will do, is a deep concern that human beings in general have. We want to know how our own life turns out, how people close to us will do, and what happens to people in general. Thus, this concern is not just about oneself: there is also about one's family, whether one's children will grow up to live a good life, and more generally whether humanity will flourish and continue to exist, or disappear and go extinct.

In general we are ignorant of the answer to all these questions. Normally one does not know how one's own life will go overall while one is still reasonably young, nor what will happen to people one is close to or to humanity itself. But if one continues to live, then one finds more and more answers to these kinds of questions. Just being alive long enough and being connected to the relevant people should give one enough information to answer basic questions like this, with the exception of the ultimate fate of humanity, which hopefully will outlast all of us. But if one dies, then one is bound to remain ignorant of the answers to these questions. If I die, then I will never know what will happen to my children and to other human beings that continue to live. My death assures my ignorance, since the facts that I want to know about are in the future and not yet settled and to be known. Death guarantees my ignorance, but in general I

would have known had I lived longer. Death thus deprives me of the knowledge I would have had. Death can deprive me of the knowledge how my own life ends if it comes suddenly in my sleep, but this can be avoided by dying with one's eyes open. But death will almost surely deprive me of the knowledge of how the life of others will go, something I am deeply concerned with. This harm generally can't be avoided.

The issue remains, though, why there is such a connection at all, in particular if we hold, as we should, that this harm is not merely based on not having one's desires satisfied. Why are these questions about my own life and that of others of deep concern to me? Why can't I just avoid these harms by caring less about my own life and the life of others? I would like to suggest that there is a connection between striving for such knowledge and the constitution of the self which makes this kind of harm unavoidable or at least a natural companion to a properly constituted self.

5 Ignorance and the self

To illustrate the connection between the epistemic harm of ignorance and the self, let's use a narrative conception of the self to bring out the idea. This narrative account of the self can be understood as a metaphysical proposal about what the self consists in, and it can be seen more metaphorically, as a way of thinking of one's own life like a story that one is the co-author, together with the world, as well as the narrator of. We don't have to focus on the metaphysical account for now, but simply using the metaphor tied to a narrative conception we can say that dying in one's sleep is a harm since one never comes to know how one's own story ends, a story one has been working on one's whole life.³ On a narrative conception of the self, the self is constituted in part by a narrative of who one is, what one values, what one hopes to be, and so on. This narrative extends back in time, including who one was, what one did value, and so on. This narrative conception can be seen as a key component of a metaphysical proposal about what constitutes the self, either in general or at least as identical over time.⁴ Or it can be seen as a useful guiding metaphor that makes use of who we are, a metaphor that is to be spelled out in a more detailed metaphysical proposal later. Since I do not have the capacity here to discuss this conception of the self in more detail, I hope that taking it as a guiding metaphor can bring out the connection between the self and ignorance, something that might well require a more serious metaphysical account of the self to defend.

How should we think of the story that in part constitutes who we are? One natural but *narrow conception* of the story of one's life is that it concerns primarily oneself and one's immediate surroundings. This story says that I got up, had breakfast, then drove my kids to school, etc.. I am the center of the story, and no one else figures except via how they are related to me. The narrow story is very egocentric, kind of a narrative solipsism. On a *wider conception* of the narrative, I am still central, but other people close to me figure as well as characters of their own right. It says that I got up, had breakfast, and then drove Freddy and Susie to school. Freddy and Susie are characters in the story just like me, although in a secondary role. It is not a solipsistic story, but one involving characters other than myself that are subjects just as I am.

³See also Velleman's own similar wording in (Velleman, 2015).

⁴See (Schechtman, 1996) and (Velleman, 2006).

On the narrow conception, the story ends with my death. It was all about me, and once I am gone there is nothing left to say about this story. But on the wider conception of the story, the main character disappears, but there is more to say: what happens to the other characters that prominently figured in it so far: Freddy, Susie, and whoever else came up prominently. On the wider conception the story is disrupted by one's death but does not end there. There are other characters still around, and they are around in one's own story.

Taking the narrative as constitutive of the self, this also leads to two conceptions of what is involved in this constitution. On the narrow conception it is just the more or less solipsistic facts about us. On the wider conception the self is constituted by a story that centrally involves other characters, who in turn are constituted by their own stories. We figure in each others' story and so collectively constitute ourselves and each other. How that might go more precisely is not so clear, of course, but the general picture might be clear enough to work with for a few more pages now.

Suppose then, for the moment at least, that the wide conception of the story that constitutes oneself is correct, as I think is plausible. We don't write our own stories one by one, we write our stories together, developing characters that figure prominently in each others stories, each of which is constituted not by a narrow story, but a wide, collective one. With this setup we can now appreciate a few key conclusions regarding ignorance, death, and the self.

We can note that on this wide conception we will likely never know how our own story ends. We might know when its main character disappears, since we might die with our eyes open, but we won't know how the rest of the story goes on: what happens to the other characters that are part of the story? How will they fare? What will happen to Freddy and Susie, and how will their fate in my story develop? I will never know, and I know that this is so, at least leaving aside unusual cases.

Furthermore, this ignorance about how our own story continues and eventually ends is due to death. Even if I die with my eyes open, my story goes on, but I will not know how. Normally I would have known simply by living longer, but death deprives me of this knowledge.

Knowledge of how my story continues is a matter of deep concern to me, in fact central for the constitution of the self via such a narrative. Knowledge of my own story so far is crucial for continuing it, and thus for constituting the self on the narrative conception of the self. To continue writing my story I cannot simply make things up as I go along. I need to base the next chapter of the story on the earlier chapters, and to do this properly I ideally need to know what was in these chapters. Ideally, I know how my life went so far to narrate into the future. Thus in constituting the self I have an inherent interest in knowledge of my story. And to continue to constitute the self into the future I need to continue to come to know how my life is going. The narrative conception of the self thus has a central epistemic component. Constituting the self is not simply a creative activity carried out in isolation of the facts, it is supposed to be based on facts and continues building on them. To do this properly one should come to know about these facts, and thus constituting the self this way is constitutively tied to aiming for knowledge.

This connection between knowledge and the self can be the basis for an explanation about why a certain kind of ignorance is a harm to a person. This harm is not simply based on someone not getting what they want, but it is based

on an obstacle to continuing to constitute one's self into the future. And since this active constitution of the self is one of one's core activities, knowledge of one's own story in the wide sense is of deep concern. Ignorance about these facts is thus a harm since it undermines the very constitution of the self. Of course, all this is only a first step towards an account, in particular since the narrative conception of the self was only taken as a guiding metaphor so far, and to spell it out is neither trivial nor does it obviously lead to the truth. Still, it does seem to me to point in the right direction and it would make sense of the intimate connection between ignorance of certain facts and the harm suffered from death.

6 Ignorance and time

Before we conclude I would like to point out two connected issues tied to temporal asymmetries concerning the past and the future.

First, we can now see that a central part of the harm of death is in effect the time-inverted case of the epistemic harm tied to a missing child. When a child is missing, then I don't know what did happen to them. And when we die we will never come to know what will happen to our children. Both cases are cases of ignorance, one about the past, the other about the future. The past ignorance case of a missing child is rare, but the future ignorance case is universal. Facing death is like facing the fact that all other people in one's life go missing, without any knowledge of what will happen to them. This can shed light on the gravity of the epistemic harm tied to ignorance that comes with death. Dealing with a missing child and the suffering and harm associated with the ignorance about what happened to them and where they are now is uncontroversially great, and we generally feel great pity for those who suffer this unfortunate fate. But death does the same, except everyone goes missing and everyone is confronted with this ignorance. To be clear, there are some differences that are worth making explicit: a child usually goes missing because something bad has happened, and that fact is not lost on those who are ignorant about what happened to a child. This does not quite hold for the future, although there surely is a lot of uncertainty about the well-being of future generations. But as I argued above, there is great epistemic harm to the parents of a missing child, not just great harm to the child that is missing. There would be epistemic harm to the parents of missing children, even if missing children would do well. And similarly, there is epistemic harm from death even if the future will go well.

Second, the ignorance view can straightforwardly explain temporal asymmetries in our concern about our own non-existence and the perceived difference in harm associated with pre-natal and post-death non-existence. Our unequal concern about our future and past non-existence can be based on an unequal epistemic harm connected with these two cases of non-existence. We know a lot more about the past than we know about the future. We know when and where we were born, but not when and where we will die. We generally know what happened to our parents in the past, but not what will happen to our children in the future. It is a puzzle in the philosophy of time why there is this asymmetry in our knowledge between the past and the future, but we do not need to solve this puzzle about time and knowledge here. We can accept it as a fact that we know more about the past than about the future even without understanding fully why this asymmetry obtains, and then use this fact to explain a feature of our unequal concern about past and future

non-existence. My past non-existence, i.e. the time before I was born, is a lesser source of epistemic harm to me, since I know a lot about the past, and what happened to the people close to me and to humanity. My future non-existence, on the other hand, is a greater source of epistemic harm: it prevents me from finding out about other people, in particular younger ones like my children, and the path of humanity. My future non-existence is the reason why I will not overcome my ignorance about matters of deep concern. It is because of my death that I will remain ignorant about the answers to these important questions about the future, answers I do have to corresponding questions about the past.

7 Conclusion

This then is an outline of an idea about how to better understand the harm of death. I proposed that there is a significant harm of death tied to ignorance. The significance of the epistemic harm tied to death was motivated by considering the epistemic suffering experienced in missing person cases and the epistemic harm that supports this suffering. I outlined an attempt to explain the occurrence of this harm via a connection between knowledge and the self, based on the narrative conception of the self and a wide conception of what story constitutes us. Whether this is only a helpful metaphor or closer to the literal truth of what constitutes the self is much beyond the scope of this essay. Similarly, what precisely the connection is between knowledge and the self, and the deprivation of knowledge, i.e. ignorance, and suffering and harm, needs to be developed more fully, but is also too substantial to be worked out here. A fuller story would provide a better explanation than I was able to give here of why merely an epistemic state of oneself has these significant consequences for suffering and harm. My main point was more modest, but still hopefully somewhat substantial. I argued that there is such a thing as epistemic suffering and there is a connection between ignorance and epistemic harm. Furthermore, I argued that there is a connection between ignorance and harm, in particular ignorance of certain future facts and the harm of death. There surely are other harms of death besides ignorance, but ignorance of the kind discussed is a distinct and especially significant aspect of the harm of death, one that is neglected by a hedonistic deprivation account and other approaches. Death thus comes with a distinctly intellectual harm.⁵

References

- Boss, Pauline (1999). *Ambiguous loss: Learning to live with unresolved grief*. Harvard University Press.
- Fricker, Miranda (2007). *Epistemic Injustice: Power and the Ethics of Knowing*. Oxford: Oxford University Press.
- Schechtman, Marya (1996). *The constitution of selves*. Cornell University Press.
- Velleman, J. David (2006). "The Self as Narrator". In: *Self to Self: Selected Essays*. Cambridge: Cambridge University Press. Chap. 9, pp. 203–223.
- (2015). "Dying". In: *Beyond Price: Essays on Birth and Death*. Openbook Publishers.

⁵For comments on earlier versions I am indebted to participants at the Analytic Philosophy Forum at the University of Tokyo, my philosophy of death class at UNC Chapel Hill, and to David Velleman.