

JUNK BELIEFS AND INTEREST-DRIVEN EPISTEMOLOGY*

Jane Friedman | jane.friedman@nyu.edu

12/2016

1 Introduction

In *Change in View*, Gilbert Harman considers the following epistemic norm,

Logical Closure Principle. One's beliefs should be "closed under logical implication." In other words there is something wrong with one's beliefs if there is a proposition logically implied by them which one does not already believe.

This sort of norm surely has some hold over at least some epistemologists. Harman rejects the norm on the following grounds though,

But the Logical Closure Principle is not right either. Many trivial things are implied by one's view which it would be worse than pointless to add to what one believes. For example, if one believes P , one's view trivially implies "either P or Q ," "either P or P ," " P and either P or R ," and so on. There is no point in cluttering one's mind with all these propositions. And, of course, there are many other similar examples.

And he then makes explicit the norm implicit in this rejection,

Clutter Avoidance. One should not clutter one's mind with trivialities.¹

*Please cite published version: Friedman, J. (2018), JUNK BELIEFS AND INTEREST-DRIVEN EPISTEMOLOGY. *Philosophy and Phenomenological Research*, 97: 568-583. <https://doi.org/10.1111/phpr.12381>.

¹See Harman (1986), p. 12 for this discussion.

Harman continues with a small discussion of some details (some of which we will touch upon later). As far as I can tell, epistemologists have been largely sympathetic to Harman's general thought here. A quick scan of the literature finds much (mostly casual) endorsement of at least some version of a clutter avoidance norm. For instance, Goldman (1986), Christensen (1994), Williamson (1998), Ryan (1999), Feldman (2000), Wallace (2001), Sainsbury (2002), DePaul (2004), MacFarlane (2004), White (2005), Field (2009), Douven (2010) and Wedgwood (2012) register support for some form of the thought that we should avoid cluttering our minds.

My plan in this paper is to draw out some consequences of endorsing Harman's thought. Some of these consequences seem to me to be fairly significant. Given this, I think it's worth drawing them out, especially in light of the extent to which these "clutter considerations" have resonated with philosophers. I am not going to spend much time defending Harman's claims about clutter. I add myself to the list of those who find them compelling, and I'll say a little bit in defence of Harman. My main aim here though is to draw out the picture of normative epistemology we get if we think that clutter is to be avoided. I think that picture is different from the one that many epistemologists take for granted.

2 Avoiding Clutter

In this section I want to flesh out the claim that we should avoid cognitive clutter. Harman claims that, "one should not clutter one's minds with trivialities". Let's start there.

We can assume that a necessary condition on cluttering our minds is acquiring new attitudes – the things that clutter our minds are attitudes. I see no special reason to deny that any type of attitude could count as a cluttering attitude in some sorts of circumstances. The focus of the discussion of clutter avoidance though has been on doxastic attitudes and in particular on belief. Given this, I'll keep my focus on belief as well, although I think the things I say could be easily extended to other doxastic or epistemic attitudes, and perhaps to other attitudes more generally.

Even if beliefs are the sorts of things that can clutter our minds, only some beliefs will count as the cluttering ones: Harman is especially con-

cerned with beliefs towards “trivialities”. This terminology is somewhat unfortunate since it easily brings to mind trivial truths or trivial logical consequences. But Harmanian trivialities needn’t be true or logical consequences of other beliefs, and trivial truths and consequences needn’t be trivialities. Harman also describes trivialities as “useless”, “pointless”, and “matters in which one has no interest”. Rather than call these matters ‘trivialities’, I’m going to call them ‘junk’.

Here are some things that I think we should say about junk. First, x is junk for a subject at a world and time. Second, I think that the right way to think of x here is as a subject matter. So, it’s subject matters that are junk for subjects at worlds and times. There’s a lot to say about subject matters, but we can avoid most of it. In general we can think of a subject matter as a set of propositions. It is typically thought that not all sets of propositions count as subject matters, but only ones with certain properties, e.g., partitions.² In order to stay close to the discussion of junk in epistemology though, let’s limit the focus to very specific subject matters – matters of whether P . These sorts of subject matters are two-element sets, or ‘whether’ questions. For instance, the subject matter of whether Rover ate the apple or whether R is just equivalent to the set $\{r, \neg r\}$. Those two propositions – Rover did eat the apple (r), Rover didn’t eat the apple ($\neg r$) – are the two possible answers to the question of whether R and comprise its subject matter.³

Subject matters in which one has no interest are junk subject matters for one. One has an interest in the matter of whether P at w, t iff one has an interest in either of p or $\neg p$ at w, t . And so one has no interest in whether P – that matter is junk for one at the relevant world and time – iff one has no interest in either of p or $\neg p$.⁴ In what’s to come I’m going to talk not only about junk subject matters but junk propositions and junk beliefs as well. A belief is junk (for S at w, t ; I won’t keep repeating this) iff it’s a belief towards a junk proposition, and a proposition p is junk iff p ’s corresponding

²See, e.g., Lewis (1988) for some discussion.

³I’m using uppercase letters as sentence variables and lowercase ones as proposition variables. In general, the sentence ‘ P ’ expresses the proposition that P or just p , and the question of whether P has as its two complete possible answers, the propositions that P/p and that not- $P/\neg p$.

⁴One’s interests in subject matters can vary in strength, and so we can assume that one has an interest in whether P so long as one has any interest at all in that matter, and that one has no interest in whether P otherwise.

‘whether’ question or subject matter (whether *P*) is junk.

What is it to have an interest in a subject matter? Roughly, it’s to have some interest or desire served by having beliefs (or knowing) about the relevant subject matter. Or to say the same thing: to have it be the case that having beliefs or knowledge on the subject matter be in one’s interest. We should take it that just finding some subject matter interesting is sufficient for it’s being in one’s interest to have beliefs on the matter; it’s not necessary though. One might find the matter of whether *P* dull but need to have information on it nonetheless (e.g., for one’s job). There whether *P* is not interesting for one but having beliefs on the matter is in one’s interest still. There’s obviously plenty more to say about these sorts of cognitive interests and what serves them, but hopefully what’s been said here is sufficient for the purposes of the discussion to come.

While there are possible subjects for whom nothing is junk, for a typical subject at a typical world and time there will be all sorts of subject matters and propositions that will count as junk. While what counts as junk varies across subjects, worlds and times, for a typical subject junk abounds.⁵ As we’ve already seen, many of the logical consequences of the things a typical subject already believes will count as junk for a that subject: long, weird, incongruous (subject-matter-wise) disjunctions where at least one disjunct is already known or at least believed; conditionals whose consequents are already believed and whose antecedents are long, weird, incongruous conjunctions or disjunctions; and so on. Again, there are possible subjects for whom these are not junk, but these are junky for the average subject. Junk propositions can be less strange as well. If *S* does not at all care and has no use for information about whether they had wagonways in Ancient Greece or whether Manchester’s economy is growing, then propositions on these subject matters too will count as junk for *S*.

⁵There is a train of thought in some parts of epistemology that might seem to get close to the result that beliefs that amount to knowledge are never junk. According to this train of thought all knowledge is good or valuable. This thought can, at least *prima facie*, get us to the thought that no knowledge is junk. That said, if one looks to those involved in these sorts of discussions about “epistemic value” one typically finds authors struggling to accommodate something very much like Harman’s thought about the disvalue of trivialities. That is, an “unrestricted” view of epistemic value – one according to which all knowledge is valuable or of interest for every subject – is often rejected or at least refined in order to accommodate the thought that many subject matters are not interesting or valuable overall for typical subjects. See Grimm (2009) for an overview and discussion of this debate.

Harman's thought is that we shouldn't clutter our minds with junk. When does believing junk clutter our minds? Should we avoid all junk or just some of it? All of it, I think; that is, I think that we should say that believing junk always clutters our minds. Let me say a little bit more. A too full space is cluttered, but I don't think that we should take Harman's thought to be the thought that we should keep our minds empty or that we should try not to believe or know things generally. Instead I think that his underlying thought is that our cognitive resources are limited, and as such can be wasted. If they can be wasted, then we should take care not to waste them. When we believe things that are of no interest, value or significance to us we waste those limited resources – we waste time, energy and storage space. We shouldn't be wasteful in these ways; we shouldn't believe junk at all. Just as useless trinkets can count as clutter or as wasting space no matter what else is in the space, for subjects like us junk beliefs always clutter our minds.

I take it that the thought that committing things to memory can use and so waste time and energy is in good standing. While perhaps slightly more controversial, the claim that cognitive storage space can be wasted is plausible as well. For instance, it's been argued that although long-term memory may well have unlimited storage space in some sense, our retrieval capacities with respect to this stored information is severely limited. We regularly need to pull items from long-term memory: everyday action relies on our making use of information we've stored there. But access to that information can be easily lost – even if the memory in some sense remains, it becomes mostly useless since it cannot be retrieved or can only be retrieved with great effort. If it takes me months to recall where the mustard is, then this information becomes useless (I'll starve before I can make that sandwich), and looking for it incredibly cognitively taxing.

One way that the retrieval of (older) information becomes inhibited is by learning or storing new information. When we try to recall something by, say, putting a query to memory, a variety of related memories are activated as well and compete for access to conscious awareness where only some can end up. In this way new memories can inhibit access to older ones – the new memories become associated with the same retrieval cues as older ones and so can inhibit access to the other memories associated with those cues. Moreover, retrieving an item from memory can make that item more likely to

win out in future retrieval competitions, further suppressing other memories and making them more difficult to access. Some have been arguing for the adaptive nature of certain kinds of regular human forgetting for exactly this reason – we won't be able to recall the things we need to when we need to recall them if we have too much information in memory.⁶

These sorts of considerations also lend support to the thought that for a typical subject all junk should be avoided. A typical subject stores a great deal of information that she wants, needs and cares about in memory. But any new information she stores has the potential to impede access to that information. When new information is also useful or interesting or important, then the expected benefits of taking it on can exceed the risks or expected costs. But when that new information is junk, then in a typical case the costs will exceed the benefits in expectation. From this perspective for a typical subject at a typical world and time, all junk should be avoided.

Harman takes the claim that junk should be avoided to be a claim largely about “explicit” belief (and not about “implicit” belief). Others have endorsed this form of the thought that junk should be avoided as well, e.g., Field (2009). It's not entirely clear how to draw the distinction between explicit and implicit belief. Field says the following, “Explicit beliefs are ones that are directly stored; one implicitly believes something when one is disposed to explicitly believe it should the question arise.”⁷ Taking Field's distinction seriously, I assume that the thought that clutter considerations be limited to explicit beliefs is motivated by the thought that having or coming to have dispositions to believe doesn't take up time or space in the same way that having or coming to have explicit beliefs does. Perhaps this is right, although without more in-depth investigation into these two sorts of belief, I don't know that we should be convinced.

Either way though, being disposed to believe junk means being disposed to waste time and storage space and it's not at all clear to me why we should want to bar a rational subject from believing junk but allow or even encourage her to be disposed to believe junk. If we shouldn't form certain kinds

⁶There is a huge body of literature on memory and retrieval and in particular the idea that learning new information can adversely impact the retrievability of old information. Anderson et al. (1994) is a nice place to start, and Bjork and Benjamin (2011) a good overview of later findings. For some discussion of the epistemological implications see Michaelian (2011).

⁷Field (2009), p. 257.

of beliefs or store certain kinds of information then why should we be disposed to form those kinds of beliefs and store that kind of information? This would amount to the thought that we should be disposed to do something that we ought not to do.

What seems right is that if we're asked whether p is the case and p is (e.g.) entailed by our evidence then there's at least some case to be made, other things equal, for the thought that we should answer 'yes', whether or not p is junk. But this isn't equivalent to the claim that we should add a belief in p to our set of beliefs. So while it may well be that "arguments from wastefulness" don't apply quite as straightforwardly to dispositions to believe, I don't think that means that we can easily claim that clutter considerations don't apply to dispositions to believe. I take the arguments to come to cover all of the relevant kinds of belief.

Overall then I think that Harmanian clutter considerations get us the thought that we should avoid believing junk. For creatures like us – finite and cognitively limited creatures – we should take care not to waste resources. Where this gets us, I think, is to the following norm on belief revision,

CA Necessarily, if p is junk for S at w, t , then S ought not believe p at w, t .

CA is not the only way to conceive of the normative impact of clutter considerations. I think that some have thought that clutter considerations should lead us to a weaker norm than this, e.g., one that said roughly that S is not required to believe junk (rather than that she's not permitted to).⁸ Some, but not all, of the arguments to come go through if we thought of the normative force of clutter considerations in this weaker way. I do think though that CA is a more natural way of thinking about how to be with respect to clutter. Harman's compelling thought isn't that you don't have to believe all that junk, but that's it's perfectly rationally permissible for you to, rather his thought is that you should avoid believing junk. That's to say that it's not rational to clutter your mind, and you shouldn't do it. But the thought that believing junk is permitted, but not required, leaves it perfectly rationally permissible to clutter your mind. Given this, I think that CA (or

⁸I think this is how MacFarlane (2004) is conceiving of the normative force of clutter considerations. See p. 11.

some very similar sort of prohibition) is the right expression of the normative force of clutter considerations.

A pressing question is how to think of the ‘ought’ in CA and how to think of CA’s normative force more generally. I think that there are two main options to consider: the first is that clutter considerations are “meta-normative” and the second is that they issue in a regular first-order norm on belief revision. To say that clutter considerations are meta-normative is roughly to say that they issue in something of a constraint on norms of belief revision – they tell us which potential norms on belief revision can be genuine norms and which cannot; CA will be crucial to making those determinations. I think this is roughly how Harman takes clutter considerations to bear on rational belief revision, and I’ll discuss this option in the next section. The other option is to take it that clutter considerations get us a first-order norm, i.e., that CA is simply a first-order norm on belief revision. I’ll discuss this option in section 3. In section 4 I’ll say more about the upshots of both options and where I think they leave normative epistemology.

3 Clutter considerations as meta-normative

Harman thinks of clutter considerations as meta-normative and something like CA as a “meta-principle” that serves to constrain principles of belief revision. The dialectic in MacFarlane (2004) seems to indicate that MacFarlane is thinking of CA (or something like it) in this way, too. I want to explore some results we get if we take this route.

We can say that if clutter considerations are meta-normative, then for some potential norm N , if N conflicts with CA (which is itself a norm), then N is not a genuine norm. And we can say that norms N and M conflict just in case there are cases in which no subject can conform to both N and M . For instance, if there are circumstances in which M says that S ought to ϕ and N says that S ought not to ϕ , then M and N conflict.

As we saw, Harman uses clutter considerations to argue that a particular sort of closure principle is not a norm on belief revision. Here’s a way to think of his reasoning there in the current framework. Take the relevant closure norm C which says (say) that if some proposition obviously follows from one’s body of belief then one ought to believe that proposition. Now

take a case in which a junk proposition p obviously follows from one's body of belief. In this case C says that one ought to believe p ; but CA says that one ought not believe p . This means that C and CA conflict and so that C is not a genuine norm.⁹

But if we think of clutter considerations as meta-normative and CA as functioning so as to knock out potential conflicting norms, then many more candidate norms on belief and believing will also turn out not to be genuine norms. To see this, let's start with the following norm schema,

Evidential Ought (EO) Necessarily, given that S is in evidential circumstances Ψ with respect to p at w, t , S ought to believe p at w, t .

Instances, of EO are potential norms. These potential norms issue "evidential" requirements. My thought here is that there is a kind of epistemic normativity that supervenes on or is grounded in our evidential circumstances. Certain kinds of evidentialists will think that all there is to epistemic normativity is evidential normativity, but I don't mean to be making that assumption here.

If we think of clutter considerations as meta-normative in the way just discussed, EO will have no *genuine instances* – none of those potential norms will be genuine norms. Here is one way to see this. The instances of EO – the potential norms – will replace ' Ψ ' with a specification of S 's evidential circumstances for some p . Let's think about the best sorts of evidential circumstances a subject can be in with respect to p . Here are some plausible features this set of circumstances will have. First, the evidence itself should be epistemically unassailable – it should be known and clearly or unproblematically so. Second, when it comes to the evidential support relation, it looks as though (deductive) entailment is the best we can do – that's the strongest degree of support a proposition can have from a body of evidence; any inductive or ampliative support is weaker than that. Third, bodies of

⁹One might argue that this isn't quite Harman's train of thought. Instead perhaps he's reasoning that the relevant sort of closure norm demands that we believe all of the relevant logical consequences and that would be badly cluttering since there are so many; a few of those consequences here and there might be just fine, but not the lot of them. Whatever Harman intended, I've already argued that we should be thinking of all junk as clutter and given this that we ought not to believe any of it. That said, it's worth pointing out that the arguments in this section can go through with only slight modifications even on this other way of thinking about how much clutter we should avoid.

evidence that are rich and diverse seem to be better than those that are not. Let's call the evidential circumstances with this feature set α . α should amount to something like the best evidential circumstances a subject can be in with respect to p (if you think some other features need to go in here, that's fine, the argument here should go through still).

Now, take an instance of EO in which ' Ψ ' is replaced with ' α '. Given what's been stipulated about α so far, there's no inconsistency with p 's (the relevant supported proposition) being junk for S at the relevant w, t – there is no bar to S's being in the best evidential circumstances towards junk. Say p is junk then. The instance of EO in which $\Psi = \alpha$, will nonetheless demand that S believe p at w, t . But this means that this instance of EO conflicts with CA and so is not a genuine norm for belief revision.

But if α cannot make for a genuine instance of EO, then it is difficult to see how any evidential circumstances can. Any other way of specifying Ψ will weaken the subject's evidential circumstances, and so it is hard to see how those weaker conditions could demand belief if the stronger ones do not.¹⁰ If α can't get us a genuine instance of EO, then no other evidential circumstances can and so EO will have no genuine instances – no potential norm of that form will be a genuine norm.

We can get to the same sort of conclusion by a different but nearby route. We've seen that junk can be entailed by a body of even impeccable evidence, but plausibly junk propositions can be supported to any degree whatever. We have no reason to think that there is some degree of support such that junk could not be supported to that degree; our interests do not discriminate against certain kinds of evidential support. If one thought that evidential requirements supervene on degree of evidential support, then the considerations in this paragraph can be used to show that if we think of clutter considerations as meta-normative in the sense at issue, then EO has no genuine instances; there are no "purely evidential" requirements at all since junk can be supported to any degree.¹¹

¹⁰For instance, were we to say that inductive support can demand belief but not deductive support, we could end up with the following sort of possible case: a subject with good inductive evidence for p at w, t_1 ought to believe p , but at w, t_2 her evidence improves so that it decisively proves p , and then at w, t_2 it is no longer the case that she ought to believe p . This is not a good result. Similar considerations will apply to other weakenings of α .

¹¹Here is one way to get junk having any evidential probability whatever (or at least any that a subject could have for a proposition). Some necessarily false propositions will be junk,

And we shouldn't stop at EO. Take the following norm schemas,

Evidential Permission (EP) Necessarily, given that S is in evidential circumstances Ψ with respect to p at w, t , S is permitted to believe p at w, t .

Evidential Justification (EJ) Necessarily, given that S is in evidential circumstances Ψ with respect to p at w, t , S is justified in believing p at w, t .

Reliabilist Permission (RP) Necessarily, given that S's p -belief at w, t is the product of a reliable belief-forming method m , S's p -belief is permissible at w, t .

Reliabilist Justification (RJ) Necessarily, given that S's p -belief at w, t is the product of a reliable belief-forming method m , S is justified in believing p at w, t .

The argument can be straightforwardly extended to EP and then EJ. The instance of EP where $\Psi = \alpha$ will be one that says that when S's evidential circumstances with respect to p are impeccable, she is permitted to believe p . But when p is junk for S, CA says that she's not permitted to believe p and so that instance of EP conflicts with CA. And again we can either argue that if evidence alone can't permit believing in the best case, then it can't permit believing at all. Or argue, as before, that we can find junky propositions getting most any degree or grade of evidential support. So EP has no genuine instances if CA is a meta-principle in the relevant sense. But I take it that if EP has no genuine instances, then neither does EJ. If evidence

e.g., long, weird, incongruous conjunctions involving contradictory conjuncts. Take one of these junky necessary falsehoods, call it j . Now take a typical S (and w, t) with a body of total evidence E . Say that G is the set of propositions graspable by S at w, t , and Ev is S's evidential probability function. Now let's assume the following plausible claim:

- (a) Given that j is junk for S at w, t , $(p \vee j)$ is also junk for S at w, t for each $p \in G$.

And we also have:

- (b) $Ev(p \vee j) = Ev(p)$, for each $p \in G$ (other than $\neg j$) for S at w, t .

But given (a) and (b) for any degree of evidential probability that a proposition in G gets, there will be junk with that same evidential probability. So if a possible subject has evidential probability x for a proposition, then there is junk that gets evidential probability x .

alone never permits believing, then it never justifies believing either. The assumption here is that if one is evidentially justified in believing p in a case, then one's believing p in that case is evidentially permissible as well (if you don't like this assumption, we can extend the original arguments to EJ as well).¹²

Extending the results to the reliabilist norms using the same strategies is straightforward as well. Quickly, plausibly a method that takes one from a body of known information to the obvious logical consequences of that information is among the most reliable methods that we can use. But that method is one that can have beliefs in junk propositions as its output as well. But then that method endorses believing junk by making those beliefs permissible or justified. So instances of RP and RJ in which m is the sort of hyper-reliable, deductive belief-forming method at issue now, won't be genuine norms if CA is a meta-principle constraining norms of belief revision. But again, if there are no genuine instances of those norms when m is hyper-reliable, it's hard to see how there could be genuine instances when m is less reliable than that.

If we follow Harman (and others) and think of clutter considerations as meta-normative and so take it that any potential norms which conflict with CA in the sense at issue cannot be genuine norms for epistemic subjects, then many of our most familiar (potential) epistemic norms are not genuine norms at all. This won't be to say that one's evidential circumstances or the reliability of one's methods are not relevant to whether one's beliefs are as they ought to be or as they're permitted to be, but that the sorts of purely evidential and reliabilist norms we've been investigating (and plausibly others), are not genuine norms for belief revision. Another way to think of some of the results here: on the picture we've been fleshing out – one that starts from fairly plausible premises – there are no purely evidential requirements to believe nor are there even any purely evidential permissions or justifications. Evidence alone doesn't demand belief, and it can't even, on its own, permit or justify belief. And we can say similar things about

¹²I also think that a purely evidential reason to believe is at risk now. The position according to which evidence provides or generates reasons to believe but can never permit or require belief seems to me very unstable. No matter how strong those evidential reasons, and even if there are no obvious conflicting reasons, evidence alone never demands belief and never even permits it. I'm not sure how to understand this sort of position.

reliabilist normativity.

While Harman doesn't draw out these consequences from the thought that clutter considerations are meta-normative, he does think that a natural upshot of his view is that central norms of belief revision combine considerations about one's interests with more traditionally epistemic sorts of considerations in issuing normative verdicts. For instance, he claims that one should only come to have or retain some belief on the matter of whether P if one is interested in whether P . For Harman, norms of belief revision are by and large "interest dependent", which is to say that the normative verdicts that they issue are in part a function of our (very local) interests. The position according to which the central norms of belief revision are conditional on our interests or hypothetical is not the only one left open by the arguments here, but it is a natural one and it allows us to maintain the normative influence of traditionally epistemic considerations. On this sort of picture we'll have norms that say roughly that if p is non-junk for S at w, t , then. . . . And the ellipsis here can be replaced with some of the more plausible instances of the sorts of norm schemas we considered.

Crucially, on this sort of account, the traditional domain of the epistemic is severely constrained: we've seen that all sorts of purely evidentialist and reliabilist potential norms are not genuine norms (and I assume we can show the same for other candidate epistemic norms). While the argument doesn't impugn a purely epistemic requirement (or permission or reason) to not believe some p , and perhaps leaves some other bits of what's typically thought of as the purely epistemic in tact, it looks as though there will be no purely epistemic requirements, permissions, justifications (etc.) to believe.¹³ If some belief one has is permissible or one is required to have or justified in having, there will be no cases in which those normative properties are purely epistemic, at least not as the epistemic is traditionally conceived of.

¹³See Nelson (2010) for the thought that all of our epistemic duties are negative, and that our requirements to believe are also a function of genuinely non-epistemic considerations (practical, moral, etc.). Nelson is worried that "positive epistemic requirements" – e.g., requirements to believe – would be too demanding for normal epistemic subjects to abide and so assuming that ought implies can, they should be rejected.

4 CA as a first-order norm on belief revision

In the last section I drew out some of the results that came with thinking of clutter considerations as meta-normative. In this section I want to see what happens if we think of CA as a first-order norm on belief revision instead. Here we might think of CA as epistemic or as non-epistemic. I think that versions of arguments to come can be made on either way of thinking about a first-order clutter-avoidance norm and it's not clear to me which way we should characterize the norm. More generally, the question of how we should tell the epistemic from the non-epistemic is obviously contested, and I don't feel particularly convinced by any attempted resolution of the issue. That said, I think that most people will want to think of CA as non-epistemic, and so I'm happy to go forward with that way of thinking about it.

We are not forced to think of clutter considerations as meta-normative or CA as having the power to knock potential norms out of contention, but there may be a reason that the route feels natural. The clutter considerations that resonated with philosophers weren't merely that we should worry about cluttering our minds and otherwise wasting cognitive resources, but that these sorts of worries had the power to constrain what might otherwise seem like good epistemic or doxastic practice. The thought wasn't just that we shouldn't clutter our minds with trivialities, but that we shouldn't clutter our minds with trivialities even if those trivialities were all true and could be known and obviously followed from other things true and known. So there seems to be an important sense in which clutter considerations already take account of traditionally epistemic ones. In a way this should be obvious: Harman is not pointing out to us that propositions that have no positive epistemic standing for us shouldn't be believed. Rather he is claiming that despite having what look like impeccable epistemic credentials, some things still shouldn't be believed. In this sense clutter considerations are *epistemically aware*: they have seen and accounted for the relevant propositions' epistemic standings.

The claim that clutter considerations are epistemically aware has, as we've just seen, implications about the relative "weightiness" of CA: clutter considerations render an "epistemic-things-considered" verdict in favour of CA. So even if some epistemic norm says that you may believe p , if CA says

that you may not, the first verdict is outweighed by the one rendered by CA. More specifically we can say that when CA issues a verdict which conflicts with the verdict issued by some epistemic norm, then the verdict CA issues *overrides* the epistemic verdict. And we can say that if normative verdict v_1 overrides normative verdicts v_i ($i \geq 2$) then a final normative verdict – v_f , which takes account of both v_1 and v_i – will be identical to v_1 .

Another way to conceive of CA's force on the picture we're considering now is to think of CA as functioning as a "side constraint". Nozick (1974) argues that some moral norms might best be conceived of as side constraints on our actions. Those norms constrain what we are (morally) permitted to do even in the face of other normative pressures. For instance, say there is a norm that says that in some case one ought to ϕ (say ϕ -ing maximizes well-being), but ϕ -ing, in this case, involves violating peoples' rights. If there is a side constraint that says that one must not violate the rights of others, then one is overall not permitted to ϕ in the relevant case (despite some other moral norm saying that one must). This is the sort of status we should think of CA as having. Even if there is some other (epistemic) norm that says that one ought to believe p , if p is junk, then one is overall not permitted to do what one is otherwise required to do. If CA is a first-order norm on belief revision, we should think of it as a side constraint in this sense then.

Thinking of CA as a side constraint leaves it open that there are genuine instances of the various epistemic norm schemas that we thought about in the last section – perhaps there are genuine norms that say that subjects ought or are permitted to believe p just given certain kinds of favourable evidential circumstances, or that a subject's beliefs are justified whenever reliably formed or even that our beliefs ought to be closed under (known) logical consequence. CA doesn't knock these out of contention now. That said, the normative verdicts these norms issue will be regularly overridden given CA's standing as a side constraint. Given this, the picture we end up with will be one according to which it will regularly be the case that rational subjects are overall not permitted to believe some junk p , no matter how favourable their epistemic circumstances. While they may be permitted or justified given just their evidential circumstances (for instance), overall they mustn't believe all sorts of things they are evidentially and epistemically permitted or justified in believing. In some ways then, thinking of CA as

a first-order norm gets us back some of the traditionally epistemic. How happy a place this tack leaves us in though is less clear. In the next section I'll say more about this, comparing it to the results from the last section.

5 Interests in epistemology

My thought has been that there are broadly two ways that we can think about the normative impact of clutter considerations in epistemology. Either we think of them as imposing a meta-normative constraint, constraining which potential norms of belief revision are genuine, or they come in in the form of a first-order norm on belief revision. In this section I want to discuss the upshots of these options further.

I argued that when clutter considerations are thought of as meta-normative, normative epistemology looks very different than we've been assuming. That tack left very little of the "purely epistemic" intact; a natural suggestion was that it left the norms of belief revision as largely hypothetical. There has been some debate about whether epistemic norms are hypothetical or instrumental in some sense. It's worth making clear the sort of hypotheticalism or instrumentalism we are left with when we think of clutter considerations as meta-normative in the sense at issue though. Typically when philosophers argue that norms of belief revision are hypothetical they argue that they bind subjects given their general interest in the truth or in knowing.¹⁴ The thought there is that we have a general aim of knowledge acquisition (or something like this) and following our evidence or using reliable methods serves that general aim, and so we should do those things. The sort of instrumentalism I've argued that we're left with on the relevant way of thinking of CA's force though is much more fine-grained: the normative appropriateness of some belief revision can be dependent upon our very local interests in particular matters or questions.

In arguing that epistemic norms are not hypothetical in this sense (and ultimately any), Thomas Kelly relies on particular examples of cases in which subjects have overwhelming evidence for p but no interest in whether P . He claims that in these cases subjects aren't relieved of their duty to believe on their evidence.¹⁵ In a variant of a case that plays a prominent role for

¹⁴See, e.g., Foley (1987), ch. 1.

¹⁵These examples come in Kelly (2003), especially pages 625-8. The discussion of these

Kelly we imagine a subject who doesn't want to know how some film ends (or whether it ends this way or that) – she has no interest in this matter. Nonetheless, Kelly contends, when her friends tell her the film's ending, she's not rationally permitted to ignore the upshot of their testimony, rather she's required to believe as they say. That's to say, according to Kelly, this disinterested subject is still rationally required to believe in accordance with her evidence for p even though p is junk for her. Many have found Kelly's verdict here intuitively compelling. But on the sort of instrumentalism we end up with when clutter considerations are thought of as meta-normative, Kelly's disinterested subject is not at all required to respond to the testimony of her friends by believing in accordance with the evidence they provide. In fact, not only is it rationally permissible for her to fail to believe in accordance with the evidence she receives, but she's rationally required to do that (and she shouldn't even come to believe the evidence). Even impeccable evidence has little independent normative force – without a local interest we shouldn't believe.

If we think of CA as a first-order norm on belief revision we get somewhat of a different result, although it's not clear how different. On this sort of picture, assuming that CA is epistemically aware or functions as a side constraint in the way I argued it would have to, we don't get the result that norms of belief revision are hypothetical, but we nonetheless get the result that overall a rational revisor regularly ignores her evidence (and other relevant bits of her epistemic circumstances), leaving us in a somewhat similar place. In fact, I think that this result is less palatable than the one according to which norms of belief revision are genuinely hypothetical and I want to briefly say why. To preview: if we think of CA as a side constraint then we'll have to say that subjects should act against standing non-hypothetical requirements largely because some relevant subject matter does not interest them.

To get at this thought in more detail and say why I think we shouldn't like it, I want to distinguish between some ways in which norms can be “interest dependent”. Let's distinguish between conditional and unconditional norms. Unconditional norms say simply that subjects ought or may (etc.) do

cases and instrumentalism in epistemology in general that ensues in Leite (2007) and then Kelly (2007) is highly relevant, I think, to the sort of instrumentalism CA leaves us with.

something or other while conditional norms say that subjects ought or may (etc.) do something or other conditional on certain conditions obtaining or conditional on those subjects being in certain circumstances, e.g., having certain kinds of evidence, being in the position to help, and so on. We can think of these conditions – c – as the conditions upon which a conditional norm’s normative verdicts depend. Now let’s say that for any conditional norm N and any c upon which N ’s verdicts depend, N is *interest sensitive* if facts about what interests the subject are in c , and *interest insensitive* otherwise. Hypothetical norms, as they are standardly understood, and as I have understood them here, are interest sensitive (conditional) norms. Let’s then say that a norm N is *interest defeasible* iff given a subject’s interests in a case, N can be overridden, and *interest indefeasible* otherwise.¹⁶ Finally, let’s say that N is *interest involving* iff it is either interest sensitive or interest defeasible.

While the sorts of epistemic norms we end up with if we take CA to be a side constraint are not interest sensitive, they are interest defeasible and so interest involving. Let’s call norms with this sort of interest profile *quasi-hypothetical* (I hope for obvious reason). I think that thinking of epistemic norms as quasi-hypothetical is conceptually coherent, but, as I said, I think if our options are to think of them as hypothetical or quasi-hypothetical, we should opt for the former. The latter’s main virtue is that it can leave us with non-hypothetical epistemic norms. Presumably on most any account our epistemic norms or their verdicts can be overridden in some cases, but on the picture according to which our epistemic norms are quasi-hypothetical they are regularly overridden. And moreover that happens when we simply have no interest in conforming. These don’t feel like very weighty considerations – not the sort that should be able to override or defeat standing, non-hypothetical epistemic requirements or permissions or justifications. Imagine a moral side constraint that said ‘don’t do anything you don’t want to’.

While there may be space for interest insensitive but interest defeasible norms, these feel strange and aren’t easy to find in other normative domains. While categorical norms are typically defined in terms of something like

¹⁶How do interests do their overriding and defeating? One way is via interest-related norms that do the work. It’s not entirely clear how the details will go in every case, but for the main focal case, we know exactly what the interested-related overriding side constraint looks like – it’s just CA.

interest insensitivity, without mention of interest indefeasibility, all of our central categorical norms are both insensitive and indefeasible – they are not interest involving. And for good reason. Part of what we think is at the centre of “categoricalness” is that the fact that one has no interest in acting some way, or no goal served by acting that way, cannot make a failure to act permissible overall. That I have no desire to save the drowning swimmer and no interest served by saving him (or even interests served by letting him drown) don’t make it overall permissible for me to let him drown. It’s not just that my obligation to save him doesn’t seem to be hypothetical, but that in a fully general sense my lack of interest in saving the swimmer doesn’t make it overall permissible to leave him to drown – whatever the structure of the relevant norms. In fact, this general thought feels like part of what it means to say that those moral verdicts or the norms from which they stem are categorical, i.e., the norms are not interest involving in our sense.

So while we can make sense of interest insensitivity without interest indefeasibility, the position feels like an odd sort of hybrid where we at once say that the norms are really not hypothetical, they really are insensitive to our interests, but that they are easily defeated or overridden when we don’t have the right sorts of interest profiles. The position is strange, it doesn’t easily find a home in other normative domains and at bottom the result that we have non-hypothetical normative verdicts that can then be overridden by something as mild as our lack of interest in some subject matter seems like a bad one to me. While we feel as though we get to retain some purely epistemic norms when we take CA to be a first-order norm on belief revision, these epistemic norms seem incredibly weak now. This is why it seems to me a mistake to treat CA as a side constraint.

I’ve tried to argue that thinking of CA as a side constraint leaves normative epistemology worse off than the option according to which CA functions to constrain which potential norms can be genuine. While the latter leaves little by way of the “purely epistemic”, the former leaves us with what seems to me to be an unpalatable picture of epistemic norms. So I think that if we find clutter considerations compelling we do best to think of norms of belief revision as hypothetical in the sense I’ve described. Either way though, I think that the domain of the “purely epistemic” has to be re-thought somewhat if we accept something like CA – it’s either virtually empty or virtually

powerless.

If one likes CA then, it's natural to think of many of the central norms of belief revision as either hypothetical or as quasi-hypothetical. In either case, what we are permitted to believe is a function of our very local interests in subject matters. And more specifically, without a local interest in whether P , believing p – no matter what one's epistemic situation – will not be (overall) permissible.

I assume that some will be fine with the thought that the norms of belief revision are dependent on our local interests in the ways at issue now. I think though that that sort of calm often comes along with thoughts that we can theorize about a purely epistemic sort of rationality without worrying about the relevant subject's interests or that we can theorize about ideal rationality without attending to the limitations of the less-than-ideal and their implications. But as we've seen, some ways of thinking about the upshot of clutter considerations don't leave us with much in the way of the purely epistemic. And it's unclear, I take it, how exactly we want to think of the ideal reasoner and even less so whether we want to think of her as having unlimited computational power, unlimited storage space, unlimited recall abilities and so on. But more importantly, on both of these points: we are cognitively limited creatures. So even if there is some way to theorize about interest-independent epistemic norms, we should also want to know how creatures like us, creatures who seem to face normative pressures that stem from our limitations, should revise their beliefs. We all have bodies of evidence that support having all sorts of beliefs, many of which we haven't formed, and in perception we are faced with far more information than we can take in or believe. I take it that we would like to know how limited creatures like us should best manage this informational deluge. Keeping cognitive clutter down seems a key element of a good information-management strategy. As I've tried to argue here though, that has some serious implications for normative epistemology.

6 Concluding remarks

In this paper I've tried to draw out some consequences of taking seriously the thought that we should avoid cluttering our minds. One natural result,

as we've seen, is a picture of rational belief revision according to which our very local interests in specific subject matters have a great deal of normative influence. If one isn't interested in some subject matter, then having beliefs on the matter is impermissible or irrational, whatever one's evidence or other epistemic circumstances. More generally, questions about whether a rational subject ought to adopt some new belief are to be answered by looking to her epistemic circumstances as well as the extent to which she is interested in the relevant topic, how that interest measures up against those in other matters she already believes or could easily come to, and more. This isn't to say though that clutter considerations push us to the conclusion that every belief revision should be answerable to these sorts of normative pressures. Plausibly, as new belief adoption or addition should be, so too belief subtraction or deletion. Keeping clutter down plausibly requires regular pruning – as new additions need to be made these can easily come along with pressures to forget or subtract other less important beliefs, and as interests shift matters that we once cared about can turn to junk. Not all rational belief transitions need to be a function of clutter considerations for all I've said though. Rationally rejecting and so not believing some p might be made on the basis of (say) evidential considerations alone; nothing I've said here closes off that possibility. Subtractions or deletions sometimes may well be as well.

We got to these conclusions by focusing on one way of interpreting an intuitive thought that wasting our limited cognitive resources is a bad thing. In particular, I've followed Harman in thinking that the right way to express this badness is via something like CA. I think this is close to the intuitive thought that has resonated with many about cognitive clutter – no matter how good your epistemic standing is with respect to p , if p is clutter, don't believe it. The rest of the paper simply drew out the natural consequences of this thought for normative epistemology. That said, one might try to capture some of the intuitive thinking behind clutter considerations in a way that will leave us with some different consequences for epistemology. For instance, as mentioned earlier, instead of CA one might try to capture clutter thoughts by way of a more modest norm that says that believing junk is permitted, but not required. This might leave some, but not all of the conclusions here intact, e.g., it can still leave us without epistemic require-

ments to believe, but might have less of an impact on epistemic permissions or justifications. Alternatively, one might try abandoning the thought that CA is epistemically aware – perhaps clutter considerations give a subject some reason not to believe or even generate a requirement not to believe, but those don't override our standing epistemic requirements and the like. As I've already argued, I don't think either is the right way to think of the situation. And of course all of the conclusions here are conditional on CA; although many have found it plausible it is surely not unassailable.

If we do like CA though and so want to take on the sort of picture of rational belief revision that one gets from taking it seriously, I've tried to show that we're left with a highly interest-driven picture of how we ought to revise our doxastic states. One upshot of the picture though is that some revisions that are epistemically impeccable in some sense – based on perfectly good and strong evidence, the product of reliable methods and epistemically virtuous in a great number of ways – might turn out to have been irrational revisions at the end of the day since they involved taking on junk opinions. But we have no reason to think that these junk opinions can't constitute knowledge in some cases; that p is junk in the relevant sense does not seem incompatible with knowing p . Some knowledge will then be junk knowledge – knowledge that we plausibly ought not to have (sometimes when we ask, 'why do you know all of that?' we are picking up on someone's junk knowledge). There is a feeling in some parts of epistemology that you can't really go wrong by knowing, but once we start to take seriously the costs of believing some things (and so knowing), that position is difficult to maintain.¹⁷ Just as we sometimes shouldn't believe despite things going perfectly well from the perspective of our evidence and the like, sometimes we also shouldn't know.¹⁸

¹⁷Of course, knowledge can also be bad for us in some more familiar ways – knowing some things can cause us pain and suffering. In these cases though, there is typically some non-epistemic state that is the source of the pain, e.g., that my team lost the match, and our knowledge is simply what connects us to that painful state of affairs.

¹⁸For extremely helpful discussion and comments thanks to Dave Chalmers and Amia Srinivasan. Thanks also to Cian Dorr, Jim Pryor and an anonymous referee. For comments on an ancestor of this paper thanks to Ralph Wedgwood and Tim Williamson.

References

- Anderson, M. C., Bjork, R. A., and Bjork, E. L. (1994). Remembering can cause forgetting: retrieval dynamics in long-term memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 20(5):1063–1087. 6
- Bjork, R. A. and Benjamin, A. (2011). On the symbiosis of remembering, forgetting, and learning. In Benjamin, A. S., editor, *Successful remembering and successful forgetting: A festschrift in honor of Robert A. Bjork*, pages 1–22. New York: Psychology Press. 6
- Christensen, D. (1994). Conservatism in Epistemology. *Noûs*, 28(1):69–89. 2
- DePaul, M. R. (2004). Truth Consequentialism, Withholding and Proportioning Belief to the Evidence. *Philosophical Issues*, 14(1):91–112. 2
- Douven, I. (2010). The Pragmatics of Belief. *Journal of Pragmatics*, 42(1):35–47. 2
- Feldman, R. (2000). The Ethics of Belief. *Philosophy and Phenomenological Research*, 60:667–695. 2
- Field, H. (2009). Epistemology without metaphysics. *Philosophical Studies*, 143(2):249–290. 2, 6
- Foley, R. (1987). *The Theory of Epistemic Rationality*. Harvard University Press. 16
- Goldman, A. I. (1986). *Epistemology and Cognition*. Cambridge, Mass.: Harvard University Press. 2
- Grimm, S. (2009). Epistemic Normativity. In Haddock, A., Millar, A., and Pritchard, D., editors, *Epistemic Value*. Oxford: Oxford University Press. 4
- Harman, G. (1986). *Change in View*. MIT Press. 1
- Kelly, T. (2003). Epistemic Rationality as Instrumental Rationality: A Critique. *Philosophy and Phenomenological Research*, 66(3):612–640. 16

- Kelly, T. (2007). Evidence and Normativity: Reply to Leite. *Philosophy and Phenomenological Research*, 75(2):465–474. 17
- Leite, A. (2007). Epistemic Instrumentalism and Reasons for Belief: A Reply to Tom Kelly’s “Epistemic Rationality as Instrumental Rationality: A Critique”. *Philosophy and Phenomenological Research*, 75(2):456–464. 17
- Lewis, D. (1988). Relevant Implication. *Theoria*, 54(3):161–174. 3
- MacFarlane, J. (2004). In What Sense (If Any) Is Logic Normative for Thought? Manuscript. 2, 7, 8
- Michaelian, K. (2011). The Epistemology of Forgetting. *Erkenntnis*, 74(3):399–424. 6
- Nelson, M. T. (2010). We Have No Positive Epistemic Duties. *Mind*, 119(473):83–102. 13
- Nozick, R. (1974). *Anarchy, State and Utopia*. New York: Basic Books. 15
- Ryan, S. (1999). The Logic of Rationality. *Philosophia*, 27:287–299. 2
- Sainsbury, M. (2002). What Logic Should We Think With? In O’Hear, A., editor, *Royal Institute of Philosophy Lecture Series*. Cambridge: Cambridge University Press. 2
- Wallace, R. J. (2001). Normativity, commitment, and instrumental reason. *Philosophers’ Imprint*, 1(3):1–26. 2
- Wedgwood, R. (2012). Justified Inference. *Synthese*, 189(2):1–23. 2
- White, R. (2005). Epistemic Permissiveness. *Philosophical Perspectives*, 19(1):445–459. 2
- Williamson, T. (1998). Conditionalizing on Knowledge. *British Journal for the Philosophy of Science*, 49(1):89–121. 2