

IN DEFENSE OF CLUTTER

BRENDAN BALCERAK JACKSON
Heinrich Heine University Düsseldorf

DAVID DIDOMENICO
Texas State University

KENJI LOTA
University of Miami

1. Introduction

Gilbert Harman's famous principle of Clutter Avoidance commands that "one should not clutter one's mind with trivialities" (1986: 12). Harman appeals to this principle in the course of his well-known argument against logical closure, the view that one ought to believe all the logical consequences of one's beliefs. Harman's rationale for the principle is that one's cognitive resources are limited, and ought to be used wisely; one ought not waste them by forming and maintaining beliefs that are in some sense trivial. Many epistemologists have been inclined to accept Harman's principle, or something like it.¹ This is significant, because the principle appears to have significant implications for our overall picture of epistemic normativity. Jane Friedman has recently argued that one potential implication is that there are no genuine purely evidential norms on belief revision—that "evidence alone doesn't demand belief, and it can't even, on its own, permit or justify belief" (2018: 576). Rather, genuine norms of belief revision must "combine considerations about one's interests with more traditional epistemic sorts

1. See, for example, Christensen (1994), DePaul (2004), Field (2009), Friedman (2018), MacFarlane (2004), Wallace (2001), Wedgwood (2012), and White (2005). Douven (2010) expresses ambivalence about the principle, on the basis of doubts about whether believing more (or more complex) propositions really is a burden on our cognitive resources.

Contact: Brendan Balcerak Jackson <bbalcerakjackson@gmail.com>
David DiDomenico <ddidomen@mdc.edu>
Kenji Lota <kenji.lota@miami.edu>

of considerations in issuing normative verdicts” (2018: 576).² Even if we insist on keeping purely evidential norms, Friedman argues, the need to avoid clutter forces us to acknowledge that the verdicts of such norms can be overridden by consideration of our interests: even if one’s evidence requires (or permits) one to believe that p in a certain situation, it might still be the case that one is in fact *not* permitted to believe that p because doing so would violate the clutter avoidance principle. Either way, Friedman argues, accepting the principle leads to a picture of epistemic normativity that is highly “interest-driven,” a picture according to which our practical interests have a significant role to play.

These are quite radical implications, in our view, and so we take the question of whether we really should accept a clutter avoidance principle to be very pressing. Is it a genuine norm on belief, as Harman and Friedman and others think? Our aim in this essay is to argue that it is not—at least, not when the central notion of clutter is understood in such a way that the principle really does threaten to lead to a picture of epistemic normativity as having a significant practical dimension.

Philosophers who have discussed clutter avoidance have tended to focus on it exclusively as a norm governing belief or judgment. But there is more to inquiry than taking a stand on whether this or that proposition is true. At the very least, inquiry also involves asking questions and considering possible answers to them, in addition to settling them by believing and judging.³ These non-doxastic stages of inquiry, and the norms governing them, have been comparatively rather unexplored. In the concluding section of the essay, we consider whether anti-clutter considerations might be more plausibly located at a non-doxastic stage of inquiry, and we suggest that if there is a genuine anti-clutter norm, it is not a norm of belief revision, but rather a norm that governs the asking and considering of questions.

Friedman’s formulation of the principle is the most explicit and fully developed one that we are aware of, and so we take her formulation as our point of departure. We present that formulation in the next section and note that it can be construed in two different ways: as what is sometimes called a *state norm* on belief, or else as what is sometimes called a *process norm*; we suggest that the principle is most plausible when construed as a process norm. In Section 3 we take up the principle on its process norm construal and argue that it is too strong

2. Nelson (2010) also argues that positive epistemic duties such as those that say one ought to believe that p are not determined solely by evidence, but also by one’s interests, needs, desires, and preferences.

3. We won’t make any detailed assumptions about the structure of inquiry here. But we assume that it at least involves somehow asking oneself a question and that it eventually leads to one settling that question by forming a belief or making a judgment. (See Friedman in press) Inquiry in this sense is distinct from the speech act of inquiring, which is plausibly governed by a different set of norms (see Whitcomb 2017).

to be plausible. In Section 4 we turn to a way of developing the idea that anti-clutter considerations concern non-doxastic stages of inquiry, by formulating a norm governing the asking and considering of questions.

2. Principle (CA)

The principle from Harman quoted above exhorts us not to believe “trivialities.” One ought not believe all the logical consequences of one’s beliefs, according to Harman, because many of them are trivialities. As Friedman emphasizes, however, the issue is not just with trivial truths or trivial logical consequences. According to Harman, one should not believe things that are pointless or useless or irrelevant to anything one cares about; this can apply even to distant, non-obvious logical consequences of one’s beliefs. The intuition underlying the clutter avoidance principle is that we ought not to waste our cognitive resources by forming beliefs about things we have no interest in, beliefs that serve no purpose for us.⁴

Friedman captures this intuition by introducing the notion of something being *junk* for a subject at a world and time. In the first instance, ‘junk’ applies to subject-matters: a given subject-matter is junk for a subject S (at world w and time t) just in case no interest of S’s (at w and t) is served by having beliefs or knowledge about that subject-matter. Friedman thinks of subject-matters as being specified by whether-questions, such as the question of whether it will rain in Miami. If this subject-matter is junk for S (at w , t) then S has no interest in the question of whether it will rain; no interest of S’s is served by settling the question either way. We can represent a whether-question as the pair of propositions corresponding to its possible answers. For example, the rain question can be represented as the pair of propositions $\{that\ it\ will\ rain, that\ it\ will\ not\ rain\}$. This lets us extend the notion of junk to propositions: a proposition p is junk for S (at w , t) just in case the question of whether p is junk for S (at w , t).⁵

4. Why would such beliefs be a waste of cognitive resources? One concern has to do with storage space: they might take up space that would be better taken up by beliefs that do serve a purpose. But storage space may well not be scarce enough to warrant much concern over pointless or trivial beliefs (at least not on any reasonable assumption about how widespread such beliefs really are). A second and perhaps more pressing concern has to do with the overall speed and accuracy of cognitive processing. For example, some memory researchers have suggested that limitations on memory retrieval are necessary for speedy and accurate information processing (Bjork 1989).

5. On the coarse-grained possible-worlds conception of propositions typically assumed in the linguistics literature on questions, the proposition $\sim\sim p$ is identical to the proposition p . This implies that p qualifies as junk for S (at w , t) just in case $\sim p$ also qualifies as junk for S (at w , t), since the question whether p (i.e., $\{p, \sim p\}$) and the question whether $\sim p$ (i.e., $\{\sim p, \sim\sim p\}$) are identical.

With the notion of junk so defined, Friedman formulates the clutter avoidance principle as follows:

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

In short, (CA) says that you are required not to believe p if no interest of yours is served by settling the question whether p one way or the other.⁶

How are we to understand the normative force of (CA)? Friedman suggests two ways. The first way is as a meta-principle that applies to first-order norms of belief revision. On this understanding, the force of (CA) is to overrule candidate norms that conflict with clutter considerations: if a candidate principle N implies that a subject S ought to or may believe p while (CA) implies that S must not believe p then N cannot be a genuine norm—that is, N cannot be a principle that really does govern our beliefs.⁷ This understanding of (CA) coheres well with Harman’s use of anti-clutter considerations to argue against the logical closure norm of belief revision. According to Friedman, accepting (CA) on this understanding is what leads to the conclusion that all genuine norms of belief revision must be interest-involving.

However, Friedman notes that we need not think of (CA) as a meta-normative principle. We could instead think of it as just another first-order norm telling us what we are required (not) to believe. On this way of construing it, (CA) is “epistemically aware,” in the sense that it has a normative force that already takes into account the epistemic standing of the propositions in question. This means that if a certain proposition is junk for you then you ought not believe it even if it has “impeccable epistemic credentials”—those credentials are outweighed or overridden by the requirement not to believe junk imposed by (CA).⁸ This first-order understanding of (CA) does allow for the existence of purely epistemic, non-interest-involving norms on belief. But according to Friedman, (CA) on this understanding implies that whatever purely epistemic norms there are will be regularly overridden by consideration of the subject’s interests.

6. Notice that Friedman’s way of formulating the principle either rules out moral reasons to believe or else (wrongly) assumes that S has a moral obligation to believe that p only if S herself has an interest in whether p . See Feldman (2000) on moral reasons for belief.

7. Starting here we drop Friedman’s relativization to worlds and times except when it is necessary.

8. Friedman suggests that an alternative way of making sense of (CA) is as a “side constraint” of the sort talked about in Nozick (1974). A side constraint imposes limits to another norm which in effect forbid one from doing what the norm requires one to do under certain conditions. If the norms say that one ought to believe p , but p is junk (at w, t), then believing p (at w, t) is forbidden. We may believe p only as long as it is within the limits imposed by the side constraint.

In order to keep the discussion simple, our focus in what follows will be on (CA) construed as a first-order norm. But the main things we have to say about (CA) on this construal carry over straightforwardly to (CA) on its meta-normative construal as well.⁹

Aside from the question about how to understand the normative force of (CA), which Friedman addresses, there is a second question of interpretation that Friedman does not explicitly address. Notice that (CA) can either be understood as what is known as a *state norm*, or else as what is known as a *process norm*.¹⁰ On the former understanding, (CA) specifies a requirement on the overall state of mind of the subject. It says that a subject for whom *p* is junk ought not to be in a state of mind that includes the belief that *p*. If she is in such a state of mind then she is open to criticism; she is “not entirely as she ought to be,” to borrow a phrase from MacFarlane (2004). On the understanding of (CA) as a process norm, by contrast, it places a constraint on what *S* is to *do*: it says that a subject for whom *p* is junk ought not to judge that *p*, where judging that *p* is understood as coming to believe that *p*. Friedman frequently refers to (CA) as a “norm of belief revision,” which suggests that she has a process norm reading in mind. Moreover, a state norm interpretation of (CA) would make normative demands that are far too strong to be plausible.¹¹ For these reasons, we will focus on (CA) as a process norm.

9. We formulate (CA) using ‘must’ rather than ‘ought’ as Friedman does. The reason is that we will emphasize in Section 4 that what one ought to do is not to be conflated with what one must do, and it is clear from Friedman’s discussion that it is really the latter notion she has in mind. We don’t think this makes a difference to the critical discussion of (CA) to follow; the problems we will raise count equally against (CA) on both interpretations. We will return briefly to this issue in Section 4.

10. The distinction is first explicitly introduced in Kolodny (2005); for further helpful discussion, see Chapter 3 of Kiesewetter (2017). The distinction arguably has its roots in Sellars’s (1956) distinction between *ought-to-be* and *ought-to-do* norms; see also Forrester (1996).

11. For example, a thinker would be subject to criticism merely for *retaining* a belief—even a very highly justified belief—just because the thinker’s interests happen to have changed. For example, suppose that at time t_1 you have a highly justified belief that a certain burner on your stove runs a little hot. At t_1 the proposition that the burner runs hot is not junk for you, because several of your interests (having to do with cooking well, preventing fires, etc.) are served by knowing whether this proposition is true. But now flash forward a few years to t_2 , a time when you live in a new home with a new stove. It is quite likely that none of the interests you now have at t_2 are at all served by retaining information about that old stove. And yet suppose that you happen to retain the information nonetheless; you still believe that the burner on that old stove ran a little hot. Intuitively, there is nothing wrong with you or your state of mind at t_2 ; we all retain pointless little bits of information from our past. Indeed, we might even praise your excellent memory. And yet according to a state norm interpretation of (CA) there *is* something wrong with your state of mind at t_2 : it is a state of mind that you should not be in, because it includes a belief in a proposition that is now junk for you.

3. (CA) as a Process Norm

Let us turn now to (CA) on its process norm construal, as a norm about what to do in the process of belief formation and revision. In fact, (CA) on this construal can be interpreted in two ways, as either a “narrow-scope” or a “wide-scope” norm:

(CAP-n) Necessarily, if p is junk for S at t then S must see to it that: S does not judge that p at t .

(CAP-w) Necessarily, S must see to it that: if p is junk for S at t then S does not judge that p at t .

The ‘must’ in (CAP-n) is detachable, in the sense that if p is junk for S then it follows from (CAP-n) that S must not judge that p . However, the ‘must’ in (CAP-w) is non-detachable: even if p is junk for S , it does not follow that S must not judge that p . What (CAP-w) requires is merely that S be such that the conditional in the scope of the ‘must’ is true of S . S can see to it that this is so by refraining from judging that p ; but she can also see to it that this is so by adjusting her interests in such a way that p is no longer junk for her.

We do not think either of these principles should be accepted—at least, not as long as the notion of junk employed in them is interpreted in a way that really supports a picture of epistemic normativity as genuinely sensitive to practical considerations. We first focus on problems for the narrow-scope (CAP-n), and then show why those problems are not avoided by switching to the wide-scope (CAP-w).

Suppose that you believe with very good evidence that it will rain. And suppose that you believe, with equally good evidence, that if it will rain, then your next-door neighbor will bring an umbrella to work. An obvious logical consequence of these two beliefs—one that we may suppose you are fully able to appreciate—is that your neighbor will bring an umbrella to work. This proposition may well be junk for you: it may be that no current interest of yours would be served by settling the question of whether your neighbor will bring an umbrella to work. But suppose that in spite of this you ask yourself the question anyway. Having asked it, isn’t it clear that you ought to answer it in the affirmative? At the very least, isn’t it clear that you *may* do so? The question has come up, and you can easily see the answer to it—it is an obvious logical consequence of two very well-supported beliefs of yours. How can you be criticized for answering it? According to (CAP-n), however, you are *not* permitted to answer it, because this would be to judge that your neighbor will bring an umbrella to work, and that proposition is junk for you right now. (CAP-n) says that you are required not to make this judgment in this situation, no matter how obvious it is to you

that the answer follows from your other beliefs, and no matter how much support you have for those beliefs. This is very hard to accept.

Notice that we are not imagining a situation in which you are cranking out all of the logical consequences of your beliefs one by one, adding them to the stock of your beliefs regardless of whether they are junk or not. It is not hard to see why this could be considered a waste of cognitive resources.¹² Rather, we are imagining a situation in which you are asking yourself whether p , and straight-away settling on p as the answer because it is an obvious logical consequence of other propositions you believe. We think it is a stretch to call this a *use* of one's cognitive resources at all, but if it is, it is surely not a *misuse*. Perhaps you can be criticized for allowing yourself to become distracted by this question, rather than focusing on some other question that would be in your interest to answer; we will return to this suggestion in the next section. But given that the question *has* come up, you are surely not criticizable for settling it in the way supported by your evidence. Thus, if the prohibition against forming junk beliefs in (CAP-n) has any normative force at all, it is not enough to outweigh the permissibility of answering the question in the way supported by your evidence.

A natural response on behalf of (CAP-n) is to suggest that if you are asking yourself the question of whether your neighbor will bring an umbrella to work, then this question must somehow be of interest to you. If so, however, then the proposition no longer counts as junk, and you do not violate (CAP-n) by answering the question in the way your evidence intuitively permits you to. In reply, we of course agree that when a question arises in one's thinking this is often because answering it would somehow serve one's interests. And of course, we agree that one can simply be curious about the answer to a given question, and so have an interest in settling it for its own sake. But we don't think that considering a question must *always* be in the service of some interest of yours. And we can stipulate that this is a case in which it isn't: perhaps another neighbor who lost her umbrella two months ago is pestering you with the question on the sidewalk, and this leads you to consider it even though you aren't the slightest bit curious as to whether that is the same umbrella, and even though you don't even care about replying to your pesky neighbor; or perhaps the question simply crosses your mind as your thoughts wander while you are taking your dog for a walk. There are all sorts of ways a question can come before one's mind that need not have anything to do with one's interests. (Maybe you overhear someone mentioning your neighbor's name on a crowded bus.) Given that the

12. According to the view we propose in Section 4, the subject ought not to waste time and cognitive resources by asking and considering the question of what all the logical consequences of her beliefs are, or by asking whether p , for every proposition p that is a logical consequence of her beliefs. If the subject ignores this recommendation and asks anyways, however, then of course she ought to answer it in the way supported by logic.

question has come before your mind, you surely cannot be criticized for answering it in the way best supported by your evidence, whether or not any of your interests would be served by doing so. This is something that (CAP-n) cannot acknowledge.

A determined defender of (CAP-n) could dig in her heels and insist that considering or asking oneself the question whether *p* *just is* a way of having an interest in settling the question whether *p*; therefore, as soon as the question whether *p* comes up, *p* automatically moves out of the category of junk for *S*, and (CAP-n) permits *S* to answer the question as she sees fit.

We don't want to insist that this is wrong; our aim here is not to haggle over the precise meaning of the term 'interest'. However, notice that construing interests in this way severs the link between interests and considerations that are distinctively *practical*. Merely asking oneself whether *p* is enough to make it no longer junk, whether or not settling the question has any bearing on the satisfaction of one's desires or the realization of one's aims. It no longer even matters whether settling the question whether *p* would help realize the aim of one's current inquiry; *p* might be entirely irrelevant to whatever it is you are currently trying to figure out. This is important. For the way (CAP-n) threatens to lead to an interest-driven epistemology is by forbidding the thinker from making judgments that would be required or permitted, from a purely epistemological viewpoint, just because those judgments happen not to serve any of the thinker's interests. Perhaps all the evidence points to *p*, for example, and yet you are prohibited from forming the belief that *p* because *p* is junk for you. But this prohibition loses its bite if all one needs to do to make it permissible to believe *p* is simply *ask oneself* whether *p*. From a practical point of view, it may be just as much a waste of time and cognitive resources as it was before to form the belief that *p*; and yet the belief is exempt from (CAP-n), on this extremely thin construal of interests, simply because one happens to be asking about it. Once we sever the notion of interests from genuinely practical considerations in this way, (CAP-n) no longer threatens to inject a genuinely practical dimension into the realm of epistemic normativity.¹³ For all (CAP-n) has to say, the norms that apply to any question you ever consider may be entirely epistemic ones, with no hint of the practical.

Readers familiar with Friedman (2020) might notice that the neighbor's umbrella case appears to have something in common with several cases dis-

13. At least one of the authors thinks that epistemic norms are in fact best understood as norms about how to settle questions that arise in one's thinking. (See Balcerak Jackson & Balcerak Jackson 2019.) If this is correct then there won't even be a *prima facie* conflict between purely epistemic norms and (CAP-n), on this thin sense of interest: epistemic norms govern how one should answer questions when they come up, and if a proposition comes up for question it cannot be junk, so it falls outside the scope of (CAP-n).

cussed there, and yet Friedman reaches a very different verdict about her cases than we do here. In one of Friedman's cases, you are in New York city, and (for some reason) you need to determine the exact number of windows on the Chrysler Building. Your time is limited, and so you decide to go to Manhattan and stand outside the building to count its windows. During the hour or so in which you are counting windows, there are many things you could rationally come to believe. For example, you see people coming in and out of Grand Central Station nearby, and so you have all kinds of perceptual evidence for beliefs about those people. You could also add to your beliefs by drawing out consequences of information you already had before arriving at the Chrysler Building; for example, you could calculate approximately how many miles you have jogged so far this year. However, if you do any of those things then you would get distracted from your task, and this may lead to an inaccurate count, or to delays as you are forced to go back and re-count. For this reason, Friedman concludes that you ought not form any of those beliefs. As in the neighbor's umbrella case, in this case you have very good evidence for believing p , and yet according to Friedman you ought not form the belief that p . Isn't this directly at odds with what we suggested above is the right thing to say about the neighbor's umbrella case?¹⁴

We don't think that it is. Friedman's case has two related features that are important for her purposes. The first is that the subject is engaged in an inquiry of some urgency: you have a question for which you need to get an accurate answer in a limited amount of time. The second feature is that the aim of this inquiry could easily be frustrated if you were to form many of the beliefs you could be forming in the situation: your inquiry requires focus and attention, and letting yourself get distracted by the people in Grand Central Station would interfere with that. These two features together are what underwrite Friedman's verdict about the case. But neither feature is required for the case of the neighbor's umbrella above. You may not be engaged in any particular inquiry at the moment the question of whether your neighbor will carry an umbrella comes up; or it may be that you are engaged in an inquiry, but one that is much less urgent. (Perhaps you are deliberating right now about whether to move out of your apartment, but your current lease doesn't expire for a few more months.)

14. If this were true, why exactly would it be a problem? There are at least two potential reasons. One is that it might suggest that cases like the neighbor's umbrella case are unlikely to be dialectically effective against Friedman herself. This is not a serious concern for us; our aim is not to refute Friedman's position, but to evaluate candidate norms of belief revision. A second reason is that it might suggest that our intuitive judgments about particular cases are too unreliable to be helpful for settling the disagreement. This is a bigger methodological issue than we can adequately discuss here. But suffice it to say that reliance on intuitive judgments about cases is part of the methodology that is widely used in this area—and indeed, it is hard to see how we could make progress on our understanding of norms of belief revision without reflecting on how candidate norms would apply to individual cases. Thanks to an anonymous referee for raising this issue.

Also, it may not be the case that the aim of your current inquiry, whatever it is, will be frustrated by forming the belief about the umbrella. This could be because the inquiry you are engaged in isn't as demanding as Friedman's of your focus and attention, and of your cognitive resources more generally. And it could also be because forming the belief about your neighbor's umbrella doesn't place very heavy demands on those resources; it is just easy for you to form the belief in the circumstances. Even in these circumstances, the belief can still be junk, and so (CAP-n) still forbids you to form it. This is what we suggest is wrong, even if Friedman is right that the demands of inquiry sometimes rule out forming a belief for which you have good evidence.¹⁵

That said, it is worth pointing out that Friedman's verdict about the Chrysler Building case is not indisputable. We agree with Friedman that there is a good sense in which you ought not to let yourself get distracted in the circumstances described. However, it is not obvious that this translates directly into the conclusion that you ought not to form the beliefs in question. Consider, for example, the judgment that more people are coming out of the station than going in. Perhaps, in the course of arriving at this judgment, you do something you ought not do. But is your mistake the act of judging itself, the act of coming to believe? A lot happens before you get to that point: you ask yourself whether more people are going in or out, you visually scan the entrance to the station, you attend to certain aspects of your perceptual experiences, and so on. These activities are already a drain on the focus and attention you need for counting the windows on the Chrysler Building, even before you get to the stage of judging that more people are coming out of the station than going in. And once you do get to that stage, the last step of actually forming the belief might hardly be any drain at all. So, it is not obvious that the right thing to criticize you for in these circumstances is forming the belief, in particular, rather than for initiating or carrying out the

15. Friedman (2019) also discusses putative cases in which one ought not to form a belief that is well supported by one's evidence. For example, suppose that you rationally judge that your lucky pencil is in the pen cup on the basis of your current visual perception. Your current perception tells you much more than just the location of your lucky pencil: it tells you about the angle the lucky pencil is leaning in the cup, what other things are in the cup, where the cup is on the desk, etc. If the aim of belief formation is to add to your knowledge as much as possible, shouldn't you have formed a much stronger belief that captured all of this (or as much of it as possible) rather than just the belief about where the lucky pencil is? However, it is important to emphasize that this case is not intended to provide support for the claim that you sometimes ought not believe what your evidence supports; as far as we can tell, there is little intuitive appeal to the suggestion that you ought not judge that your lucky pencil is in the pen cup in this case. Rather, what the case is intended to bring out is a tension between a teleological, knowledge-maximizing conception of the basis of epistemic normativity, on the one hand, and specific widely accepted norms permitting the formation of beliefs supported by evidence, on the other. (If such permissions are incompatible with a knowledge-maximizing conception of epistemic normativity, as Friedman suggests, we say so much the worse for that conception.)

process that eventually led to the formation of belief. This illustrates the point we made in the introduction, about the importance of not losing sight of the non-doxastic stages of inquiry when thinking about norms and normative judgments in this area.

What happens if we move from the narrow-scope (CAP-n) to the wide-scope (CAP-w)? Things do not substantially change. Imagine once again that you believe that it will rain, and you believe that if so then your neighbor will bring an umbrella. Unlike (CAP-n), (CAP-w) does permit you to judge that your neighbor will bring an umbrella; but it allows you to do this only if you also see to it that your interests are adjusted in such a way that the proposition is no longer junk for you. This requirement is intuitively too strong, for the same reason as before: you are surely *permitted* to answer the question in the way your evidence indicates, regardless of what your interests are. Moreover, (CAP-w) brings with it a further oddity. The situation as described is one in which the only way to respect your evidence, while also respecting (CAP-w), is to adjust your interests so as to make p no longer qualify as junk. Thus (CAP-w) seems to imply that it would be appropriate for you, *qua* inquirer, to modify your interests so as to respect your evidence. But this is implausible. In any ordinary sense of interest, it may not be at all in your interest to determine whether your neighbor will bring an umbrella to work. The mere fact that you have good evidence for p need not do anything to change that. Thus (CAP-w) sanctions what would appear to be arbitrary modifications of your preferences in this sort of scenario.¹⁶

As before, one can mitigate the force of these worries by insisting on a highly attenuated conception of interests, such that merely asking oneself the question whether p is enough to move p out of the category of junk. But as before, (CAP-w) so construed is not a serious threat for those who hope to preserve a domain of purely epistemic normativity.

Friedman's principle is just one way of explicitly spelling out Harman's anti-clutter considerations as a norm on belief. It is possible that other ways could be found that would not be vulnerable to the sorts of objections we have raised here. But this does not seem very likely to us. What our objection to (CA) as a process norm brings out is that when we are in a good rational position to settle a question under consideration, it is implausible that we are criticizable *qua* inquirers for doing so merely because that answer is uninteresting. It is hard to see how to formulate a plausible norm on belief that is able to accommodate this intuition while retaining the spirit of the clutter avoidance principle.

16. This last point is related to a general concern about wide-scope norms of rationality raised by Kolodny (2005).

4. Clutter Avoidance as a Norm on Non-Doxastic Stages of Inquiry

If there is no plausible anti-clutter norm on belief, does this mean that we should simply ignore considerations of clutter and cognitive economy entirely? We don't think so. We noted in the introduction that the formation and maintenance of belief is just one aspect of the process of inquiry—albeit the one that tends to receive the lion's share of philosophers' attention. In this section, we propose that, if there is an anti-clutter norm governing some aspect of inquiry, it does not govern the *answering* of questions, by the formation of beliefs, but rather the *asking* and *considering* of questions. Asking a question and considering possible answers to it are stages of inquiry that occur prior to settling on an answer.¹⁷ It seems plausible that inquirers ought not to waste time and energy asking and considering junk questions, questions none of whose answers would serve any interests of theirs. In this section we propose and discuss a clutter avoidance principle in this spirit.

Here is the principle we propose:

(CAAC) Necessarily, if the question whether p is junk for S at t , then S ought to see to it that: S does not ask or consider whether p at t .

(CAAC) is a narrow scope process norm: it says what an inquirer ought to do—or more precisely, it says something that she ought *not* do—when the question whether p is junk for her. (CAAC) is plausible, and the objections we raised to (CA) as a process norm in Section 3 do not threaten it. Moreover, it is at least as well supported as (CA) by considerations about the importance of conserving one's limited cognitive resources. It is plausible that we are limited in how many questions we can attend to and consider at any given time; and one is criticizable *qua* inquirer for expending those limited resources on questions that are junk.¹⁸

A core motivation for endorsing (CAAC) is that it captures our intuitions about clutter considerations without the implausible consequences of (CAP-n) or (CAP-w). Consider the case of the neighbor's umbrella. Above, we saw that it is implausible that you are required not to judge that your neighbor will bring an umbrella to work in the circumstances described. At the same time, we noted that it is plausible that you are criticizable *qua* inquirer for letting yourself

17. We adopt broad conceptions of what it is to ask and consider questions. *Asking a question* can be performing an unvocalized (internal) speech act, and it can also be coming to consciously wonder whether p . *Considering a question* is (at least) a matter of raising candidate answers to salience, checking whether one already has evidence that speaks for or against one of them, and forming strategies for answering it.

18. This proposal is not to be equated with the suggestion in Perez Carballo (2018) that some questions are more valuable than others according to a purely epistemic conception of value.

get distracted by this question in the first place. (CAAC) delivers these results straightforwardly: since the question of whether your neighbor will bring an umbrella is junk for you, (CAAC) says you should not ask or consider it; once you do, however, (CAAC) says nothing against settling it in the way supported by your evidence.

A similar explanation is available in the Chrysler Building case. If (CAAC) is true, the subject in that case is criticizable *qua* inquirer, not for forming beliefs in propositions that are uninteresting, but for asking and considering uninteresting questions that distract her from answering the central question of her inquiry, such as whether more people are going into or coming out of the subway stop. In fact, there are really *two* normative failures in the Chrysler Building case, which can be teased apart by considering a variant of the case. Suppose the subject interrupts her window counting to ask herself questions that *do* serve some interest of hers yet are still distracting. Here, too, she intuitively does something she ought not to do, even though she is not violating any norms against asking junk questions or forming junk beliefs. This brings out an important difference between two aspects of good inquiry: discernment and timeliness. The discerning inquirer asks questions that are worth asking. The timely inquirer asks the right questions in the right circumstances. It is plausible that whereas the imagined subject in the original Chrysler Building case is neither discerning nor timely, the subject in the variant case, discerning though she is, fails to be timely. For our purposes, the important point is simply that (CAAC) is intended to capture intuitions about discernment rather than timeliness; this is true of candidate clutter norms more generally.

To further appreciate the explanatory significance of (CAAC) it's worth thinking through a new case. Imagine that there is a large sum of money hidden somewhere close to you for the taking and you know that instructions for finding the money have been strategically placed in the only copy of Tekiela's *Bird Trivia: Funny, Strange and Incredible Facts about North American Birds* at your local library. Suppose that you need to actually read the book carefully in order to determine the instructions—they are embedded in the text in such a way that merely glancing through the pages won't let you know where the money is. In this way, this case is interestingly unlike the earlier Chrysler Building case—comprehending uninteresting propositions is a necessary aspect of successful inquiry.

What does (CAAC) say about this case? As you read through the book, it may well happen that you come to believe some of the propositions it contains about birds, propositions that are junk for you. According to (CAP-n), this is impermissible, but it is not impermissible according to (CAAC). And this seems plausible: even if you are not at all interested in birds, you don't do anything impermissible by retaining the information you happen to acquire while searching for the instructions. (CAAC) implies that you are only criticizable *qua* inquirer if you allow yourself to become distracted from the inquiry by asking or consider-

ing junk questions about birds. For example, suppose you start asking yourself whether female barnacle geese lay, on average, more eggs than northern shovelers each year even though having an answer to this question would not serve any of your interests. According to (CAAC) you ought not to engage in this distracting inquiry. This seems right.

The birds case is one in which it is permissible to form junk beliefs in the course of pursuing an important inquiry. One might wonder whether there are also cases where one is permitted to or even ought to ask a junk question in the course of pursuing an important inquiry. Plausibly, there are. For example, suppose that instead of a book about birds, the instructions for finding the cash are hidden in a treasure box that is being guarded by a giant troll, and the troll will only allow you to view the contents of the box if you ask and consider the junk question above about the egg-laying habits of barnacle geese and northern shovelers. Intuitively, there is a good sense in which you ought to ask and consider the question, even though answering it would not serve any of your interests. (The troll just wants you to ask it; he doesn't care whether or how you answer it.) This intuition appears to be in tension with (CAAC). However, notice that the case is analogous to other familiar cases where a subject seems to be in possession of a reason of the "wrong kind," such as cases in which the subject is offered a million dollars to believe that p , or is threatened with a bullet to the head if he doesn't.¹⁹ We're not sure what to say about such cases, but we don't think they pose any new or distinctive problem for (CAAC) in particular. They are symptoms of a more general unclarity about how to fit rational normativity into the broader normative picture.

We conclude this section with two brief remarks on the normative force of (CAAC). First, note that (CAAC) is stronger than a principle that says merely that a subject is *permitted* not to ask or consider questions that are junk. This principle is entailed by (CAAC), and it is surely true. After all, it is plausible that we are even permitted not to ask some questions that *do* serve our interests at a given time; to require us to ask and consider all the many, many questions that serve some current interest of ours would be much too demanding. Surely, then, we are also permitted to set aside some questions that do not serve any of our current interests. But we think that anti-clutter considerations support the stronger claim that one ought not to ask such questions.

Second, however, notice that to say that an inquirer *ought* to do something, in the ordinary sense of the term, is not to say that she is *required* to do it or that she *must* do it. To illustrate, suppose a stranger on the subway platform asks you which train line to take to reach a certain stop. You know that both the Red line and the Brown line usually go to that stop, but that the Brown line currently

19. See, for example, Parfit (2001) and Piller (2001), as well as Hieronymi (2005).

ends too early because of construction further down the track. Clearly, what you should tell the stranger is that he has to take the Red line, that is, that this is the line he *must* take to reach his stop. To say merely that he *ought* to take the Red line would be too weak; it could leave him with the impression that he has some other option besides the Red line (although perhaps one that is inferior, e.g., because it is much slower). Or suppose you are to attend a certain meeting that starts at 10am, but you are told that you ought to be there ten minutes early. The organizers of the meeting cannot accuse you of being late if you arrive exactly at 10am: even though you arrived later than you ought to have arrived, you were there by the time you were required to be. Still, they can criticize you: you did not do what you should have done.²⁰ Similarly, according to (CAAC) when S asks and considers a question that is junk for her, S does something she ought not do, and she can be criticized *qua* rational inquirer for doing so, even though it may not be something she is strictly forbidden from doing.

Should we accept the stronger principle that *requires* the subject not to ask or consider junk questions? Such a principle would have the force of an imperative: *If the question whether p is junk for you, don't ask it! Don't consider it! That is forbidden!* By contrast, (CAAC) has the force of a recommendation. We are not convinced that the stronger principle is correct. As far as we can see, a strict requirement against asking junk questions is stronger than what is needed to make sense of our intuitions about the cases we have considered here. To put it more evocatively, if you were employed as someone whose job is to inquire, asking and considering junk questions would clearly earn you a negative performance evaluation; but we are not convinced that it would be a firing offense. Of course, it may turn out that there are other cases that do suggest that we are forbidden from asking junk questions, or that such a principle is supported by further considerations that we have not taken into account here. If so, that is still compatible with (CAAC). In fact, (CAAC) is entailed by a principle requiring the subject not to ask junk questions, since one ought to do what one is required to do.²¹

20. These examples are adapted from McNamara (1996), an early source of the observation that 'ought' and 'must' need to be distinguished in this way. See also Jones and Pörn (1986).

21. Does a revision of Friedman's (CA) in terms of 'ought' rather than 'must' avoid the problems raised in Section 3? We don't think so. Our umbrella case seems just as much like a case in which it is not true that you ought not judge that your neighbor will bring an umbrella to work as it is a case in which it is not true that you are required not to judge that your neighbor will bring an umbrella to work. But even for those who do have the intuition that you ought not form the relevant belief in this case, this intuition is better explained in terms of (CAAC) than in terms of a revised version of (CAP-n). Imagine that a nearby friend could see that you are about to raise the question of whether your neighbor will bring an umbrella to work. (CAAC) explains the propriety of your friend interrupting you by saying, "I recommend not asking that question, since nobody really cares anyways". But suppose you don't heed your friend's advice. You ask the question and can easily see the correct answer. At this point, it would not be appropriate for your friend to say, "I recommend that you not form the belief. It would not be advisable to believe that your neighbor

5. Conclusion

If we accept (CAAC), does this lead to an interest-driven epistemology? We don't think so. (CAAC) imposes a constraint on what questions we ought to ask ourselves and spend time considering, but it says nothing about how we ought to answer them. It doesn't say anything at all about what belief state we ought to be in, or about what judgments we ought to make. As far as (CAAC) is concerned, then, the norms governing belief might be purely epistemic ones that are completely indifferent to the practical interests of the thinker.

(CAAC) is obviously not intended to be the whole story about the norms governing the asking and considering of questions. The complete normative story may also include, for example, an enkaptic norm that requires one to take the means that are necessary for answering a question one aims to answer, as suggested in Friedman (2020). There may also be norms that concern the interaction between questions and beliefs, such as a norm against inquiring whether p when one already believes that p (as proposed by Friedman 2019 and Millson 2020). The present suggestion is just that (CAAC) should be considered as a part of that overall normative story, whatever it ultimately looks like.²²

What is most important, for present purposes, is that (CAAC) illustrates a potential avenue for taking seriously the considerations of cognitive economy that Harman, Friedman, and many others find compelling, without being inevitably led to an interest-driven epistemology.

Acknowledgments

For discussion and detailed comments on earlier drafts of this paper, thanks go out to Adam Shmidt and Magdalena Balcerak Jackson. We would also like to thank an anonymous referee for *Philosophical Studies* and two anonymous referees for this journal.

will bring an umbrella to work." By locating the source of the mistake in the asking of the question that the subject has no interest in answering, (CAAC) provides us with a better explanation of the normative contours of the neighbor's umbrella case.

22. The ultimate truth or falsity of (CAAC) is likely to depend on what the aims and constitutive norms of inquiry turn out to be. In particular, (CAAC) fits well with the idea that one of the aims of inquiry is the satisfaction of one's interests. But it is not clear how well (CAAC) fits with a view according to which inquiry aims only at knowledge (or justified belief, or understanding); see Kelp (2021) for discussion.

References

- Balcerak Jackson, Magdalena and Brendan Balcerak Jackson (2019). With Power Comes Responsibility: Cognitive Capacities and Rational Requirements. In Balcerak Jackson, Magdalena and Brendan Balcerak Jackson (Eds.), *Reasoning: New Essays on Theoretical and Practical Thinking* (178–202). Oxford University Press.
- Bjork, R. A. (1989). Retrieval Inhibition as an Adaptive Mechanism in Human Memory. In H. L. Roediger and F. I. M. Craik (Eds.), *Varieties of Memory and Consciousness* (309–30). Erlbaum.
- Christensen, David (1994). Conservatism in Epistemology. *Noûs*, 28(1), 69–89.
- DePaul, M. R. (2004). Truth Consequentialism, Withholding and Proportioning Belief to the Evidence. *Philosophical Issues*, 14(1), 91–112.
- Douven, Igor (2010). The Pragmatics of Belief. *Journal of Pragmatics*, 42(1), 35–47.
- Feldman, Richard (2000). The Ethics of Belief. *Philosophy and Phenomenological Research*, 60(3), 667–95.
- Field, Hartry (2009). What Is the Normative Role of Logic? *Aristotelian Society Supplementary Volume*, 83(1), 251–68.
- Forrester, James (1996). *Being Good & Being Logical*. M. E. Sharpe.
- Friedman, Jane (2018). Junk Beliefs and Interest-Driven Epistemology. *Philosophy and Phenomenological Research*, 97(3), 568–83.
- Friedman, Jane (2019). Teleological Epistemology. *Philosophical Studies*, 176(3), 673–91.
- Friedman, Jane (2020). The Epistemic and the Zetetic. *The Philosophical Review*, 129(4), 501–36.
- Friedman, Jane (in press). Zetetic Epistemology. In B. Reed and A. K. Flowerree (Eds.), *Towards an Expansive Epistemology: Norms, Action, and the Social Sphere*. Routledge.
- Harman, Gilbert (1986). *Change in View*. MIT Press.
- Hieronymi, Pamela (2005). The Wrong Kind of Reason. *Journal of Philosophy*, 102(9), 437–57.
- Jones, Andrew J. I. and Ingmar Pörn (1986). ‘Ought’ and ‘Must’. *Synthese*, 66(1), 89–93.
- Kiesewetter, Benjamin (2017). *The Normativity of Rationality*. Oxford University Press.
- Kelp, Christoph (2021). Theory of Inquiry. *Philosophy and Phenomenological Research*, 103(2), 359–84.
- Kolodny, Niko (2005). Why Be Rational? *Mind*, 114(455), 509–63.
- MacFarlane, John (2004). In What Sense (If Any) Is Logic Normative for Thought? Unpublished manuscript.
- McNamara, Paul (1996). Must I Do What I Ought (or Will the Least I Can Do Do)? In M. Brown and J. Carmo (Eds.), *Deontic Logic, Agency and Normative Systems* (154–73). Springer.
- Millson, J. A. (2020). Seeking Confirmation: A Puzzle for Norms of Inquiry. *Analysis*, 80(4), 683–93.
- Nelson, Mark T. (2010). We Have No Positive Epistemic Duties. *Mind*, 119(473), 83–102.
- Nozick, Robert (1974). *Anarchy, State and Utopia*. Basic Books.
- Parfit, Derek (2001). Rationality and Reasons. In D. Egonsson, J. Josefsson, B. Petersson, and T. Rønnow-Rasmussen (Eds.), *Exploring Practical Philosophy: From Action to Values* (17–39). Ashgate.
- Pérez Carballo, Alejandro (2018). Good Questions. In K. Ahlstrom-Vij and J. Dunn (Eds.), *Epistemic Consequentialism* (123–45). Oxford University Press.

- Piller, Christian (2001). Normative Practical Reasoning. *Aristotelian Society Supplementary Volume*, 75(1), 175–216.
- Sellars, Wilfrid (1956). Empiricism and the Philosophy of Mind. *Minnesota Studies in the Philosophy of Science*, 1(19), 253–329.
- Wallace, R. J. (2001). Normativity, Commitment, and Instrumental Reason. *Philosophers' Imprint*, 1(3), 1–26.
- Wedgwood, Ralph (2012). Justified Inference. *Synthese*, 189(2), 1–23.
- Whitcomb, Dennis (2017). One Kind of Asking. *Philosophical Quarterly*, 67, 148–68.
- White, Roger (2005). Epistemic Permissiveness. *Philosophical Perspectives*, 19(1), 445–59.