

**Note:** This is a work in progress. Please do not cite without permission

## **Salience-Based Hinge Epistemology**

Recent work in the philosophy of attention has focused on the normative significance of attention (Wu forthcoming), and particularly of salience structures (Archer 2022). One exciting direction for exploration in this area is in the potential for developing a salience-based epistemology; i.e., an approach to epistemology that centers salience structures, rather than beliefs or knowledge, as the locus of epistemic evaluation. In this paper, I point to an important limitation in the scope of epistemic evaluation that is often missed in traditional epistemology and show how a salience-based epistemology is particularly useful for analyzing this limitation. The limitation in question, identified in recent work in hinge epistemology, concerns the necessarily *local* nature of epistemic evaluation; all epistemic evaluation must take place relative to commitments that are themselves beyond direct evaluation in the ordinary sense, in terms of justification and evidence. A salience-based epistemology provides a useful model for analyzing in detail how these hinge commitments enable epistemic evaluation: a salience-based epistemology allows us to explain hinge commitments as determinants of what information is salient, versus what information will remain part of the unquestioned background to our inquiries.

### 1. Introduction: Traditional Epistemology and Salience-Based Epistemology

Traditional analytic epistemology, though encompassing a massive literature and a wide range of topics, is animated by a relatively small set of core concerns. In particular, traditional epistemology has, to a large extent, been concerned with explaining the possibility of knowledge of an external world. Indeed, the focus on (i) the nature of perception and the epistemic status of perceptual experience; (ii) the analysis of knowledge, especially in a way sensitive to counterexamples derived from the possibility of skeptical hypotheses; and (iii) analysis of sources and limitations to epistemic justification for belief. These topics, which make up much of the content of the Stanford Encyclopedia Entry for “Epistemology”, all fit the picture of epistemology as engaging in the project of explaining how individual knowers can come to know or have justified true beliefs about an external reality (Steup and Neta 2020).

For those interested to engage in epistemic questions beyond the usual concern (or obsession) with radical skepticism, the tools and concepts of traditional epistemology—including analyses of knowledge, justification, warrant, defeaters, and so on—can be useful in some respects, but also limiting in others. For instance, the literature on epistemic injustice considers how the transmission of warrant (through testimony)—itself a focal point in traditional epistemology—can be blocked or undermined due to prejudice (Fricker 2007). Nevertheless, traditional epistemology is limited in how it can account for what might be epistemically problematic (not just morally problematic) with prejudice itself. This is because traditional epistemology is largely confined to the project of evaluating the rationality of particular beliefs, while prejudices can be epistemically distorting even when they don’t involve any false or unwarranted beliefs.

To illustrate this point, Jessie Munton provides the example of Margaret, who “holds a range of true and rationally formed beliefs about Muslims, some of which are positively valenced

and some of which are negatively valenced”, yet “Margaret almost never calls to mind the positively-valenced beliefs”, while “she frequently calls to mind the negatively-valenced beliefs” (2021, 7). Intuitively, Margaret exhibits a prejudice against Muslims. This prejudice, in addition to being morally problematic, also intuitively seems to be epistemically problematic. After all, Margaret’s systematic prioritization of negatively-valenced information over positively-valenced information seems epistemically irrational—we should not expect one sort of information to be by default more relevant to various cognitive projects than the other. Yet, by hypothesis, Margaret largely holds true and rationally held beliefs about them, and for this reason the tools of traditional epistemology struggle to explain how Margaret’s prejudice could be epistemically irrational.

Salience-based epistemology shifts the locus of epistemic evaluation, so that instead of focusing on what makes individual beliefs rationally or irrationally held, we focus on what makes a salience structure an epistemically good or bad one. I discuss some recent work on salience structures in section 3. The rough idea is that a salience structure is a way of organizing information, such that some information is more accessible (salient) from a given starting point than other information. Intuitively, some ways of structuring information are epistemically better than others given a cognitive project, even when all the same information is available in each structure. Salience-based epistemology encompasses the categories of traditional epistemology, since beliefs can be thought of as mental information storage units and hence as potential elements of salience structures. But salience-based epistemology goes beyond traditional epistemology, in applying epistemic evaluation not just to beliefs but also to the salience structures that include them.

An advantage of a salience-based epistemology is that it enables us to explain a fuller range of epistemic evaluations than traditional epistemology affords. For instance, salience-based epistemology provides a plausible way us to identify what is epistemically problematic with prejudice (Munton 2021; [ms., chapter 6](#)): prejudicial salience structures will tend to prioritize negatively-valenced information about a group over positively-valenced information even when the negatively-valenced information is not more relevant to one’s epistemic projects than the positively-valenced information. While one might attempt a similar explanation by positing a cognitive component to prejudice (e.g., such that prejudice entails having certain *beliefs* about a group), the salience-based approach is useful especially when considering instances of implicit bias, where an explicit prejudicial belief often appears to be lacking.

The analysis of prejudice is one important example of how a salience-based epistemology can help to fill in some of the explanatory gaps left by the standard framework in traditional epistemology. In this paper, I will examine another limitation in the framework of traditional epistemology and show how a salience-based epistemology can again provide a useful alternative framework. In brief, the limitation concerns the scope of epistemic evaluation. Traditional epistemology assumes that any belief is in principle subject to evaluation in terms of its rational grounding, and that if a belief can be shown to lack rational grounds it must be epistemically illegitimate to hold. This approach misses the crucial role—identified in recent work on hinge epistemology—that background certainties (‘hinge commitments’) play in enabling the sort of epistemic evaluation of interest to traditional epistemology (Moyal-Sharrock 2004; Coliva 2015; Pritchard 2016). I argue that fully taking the measure of this limitation also reveals some important constraints on the epistemic norms that can apply to salience structures. Thus, though I advocate for salience-based epistemology as a useful

alternative to the traditional epistemological framework, this paper will not remain neutral with regard to the proper development of salience-based epistemology.

In the next section, I discuss in more detail the role that background certainties play in rational evaluation, highlighting a limitation of traditional epistemology. I also identify one area where the literature on these certainties is crucially underdeveloped; namely in the epistemic relation between hinge commitments and ordinary (i.e., evidentially evaluable) beliefs. Section 3 describes the basic framework of salience-based epistemology, which analyzes how salience structures can be assessed in their ability to promote or hinder core epistemic goals. In section 4, I show how a salience-based epistemology is especially well-suited to explaining how background certainties relate to ordinary epistemically-evaluable beliefs. The core idea is that in my view hinge commitments, despite being held (or rather, *because* they are held) to an optimal degree of certainty, are remarkably *low* in salience. It is crucial to hinges that they themselves are typically not the focus of inquiry. Indeed, attempting to subject a hinge commitment to critical scrutiny is apt to lead to what Duncan Pritchard calls ‘epistemic vertigo’ (2016), the uncomfortable experience brought on by the explicit recognition that one’s ordinary beliefs, though (usually) epistemically legitimate, depend upon certain commitments that cannot themselves be rationally supported. Finally, in section 5, I discuss some of the implications of hinge commitments for salience-based epistemology. Here, I argue that the way hinges define the contours of salience maps while themselves being insulated from certain kinds of doubt rules out the applicability of certain objective and universal epistemic norms on salience structures.

## 2. Where Reasons Give Out

Traditional epistemology typically construes all belief as in principle subject to epistemic evaluation, often in terms of evidence or basis in a reliable belief-forming process. For instance, according to the prominent evidentialist approach in contemporary epistemology, the epistemic justification for an individual’s belief is always determined by the individual’s *evidence* for that belief (Conee and Feldman 2004). Epistemic theories differ in how they view the source and strength of such justification, but the shared framework is one in which if a belief is shown to lack rational grounding, then it is irrational to continue to hold that belief. Following Pritchard (2016), we can call the idea that all belief must be rationally supported to be legitimate the Universality of Rational Evaluation Thesis, or just the *Universality Thesis*.

Famously—or rather, infamously—the Universality Thesis seems to generate skeptical problems, at least in combination with some other *prima facie* plausible epistemic principles. One recipe for generating skeptical problems comes from combining the Universality Thesis with the observation that having evidence or rational support for most of our ordinary beliefs is incompatible with the truth of certain skeptical hypotheses, hypotheses for which we are in principle unable to acquire evidence that would rule them out. The most famous of these in contemporary work on skepticism is the brain-in-a-vat (BIV) hypothesis, the hypothesis that one is currently a disembodied brain suspended in a vat of nutrients and electrochemically stimulated by a supercomputer such that one has all the sensory experiences one is currently undergoing. The point is that (i) if this hypothesis were true, it would be false that I am now currently working on this paper at my desk, and almost everything else I take myself to know would be false as well, and (ii) in principle there seems to be no conclusive evidence I can have or acquire which would rule out the BIV hypothesis, since any sensory experiences

purportedly providing such evidence could instead be concocted by the supercomputer I am wired up to. Given that my having knowledge that I am currently writing at my desk (and many other ordinary pieces of purported knowledge) requires having evidence that rules out hypotheses incompatible with that knowledge, I would seem to be unable to have very much knowledge at all.

One influential way of precisifying the force of such skeptical arguments in traditional epistemology relies upon the combination of the Universality Thesis with plausible closure principles on knowledge. Closure principles provide an epistemic bridge between our ordinary claims to knowledge and the denials of skeptical hypotheses. While it is not trivial to adequately specify a closure principle on knowledge, the intuitive idea is that knowledge is closed under known deduction:

*Closure (simple):* If  $S$  knows that  $p$ , and  $S$  knows that  $p$  entails  $q$ , then  $S$  can come to know that  $q$ .

Closely related to closure principles are warrant transmission principles, which describe conditions under which epistemic warrant for some belief can be transferred to the known entailments of that belief (see Baumann 2011):

*Transmission:* If  $S$  warrantably believes that  $p$ , and  $S$  knows that  $p$  entails  $q$ , and competently deduces from belief that  $p$  and knowledge of the entailment relation that  $q$ ,  $S$  can come to warrantably believe that  $q$ .

What is important about the closure and transmission principles is that they describe a central way in which our knowledge or our warranted beliefs form a *system*. Typically, one does not have or acquire beliefs in isolation, but one does so in relation to a broader system of background beliefs. Skeptical arguments exploit these systems by introducing hypothesis which (i) one apparently cannot know or warrantably believe to be false, yet (ii) bear entailment relations to many of the other propositions one would ordinarily take oneself to know. So, taking “I have hands” as an instance of ordinary knowledge, we have the standard skeptical problem (as described in Pritchard 2016):

- I.  $S$  knows that  $S$  has hands.
- II. If  $S$  knows that  $S$  has hands, then  $S$  knows that  $S$  is not a (handless) brain-in-a-vat.

(II) reflects an instance of the closure principle. Yet, in this instance, the closure principle links an ordinary piece of knowledge to something it seems that one could not be in a position to know. That is:

- III.  $S$  does not know that  $S$  is not a (handless) brain-in-a-vat.

We seem stuck in a trilemma: either reject the closure principle, claim that we can indeed know the denials of skeptical hypotheses, or concede to the skeptic that we cannot even know that we have hands. While closure principles have not been without detractors, they are generally regarded as highly plausible (when suitably construed). Those who reject the closure principle, for instance, seem forced to endorse what Keith DeRose has aptly labeled ‘abominable conjunctions’, e.g., “I know that I have hands but I don’t know that I am not a handless brain-in-a-vat” (DeRose 1995). These conjunctions seem absurd in much the same way that Moore-paradoxical sentences do, e.g. “It’s raining, but I don’t believe that it is”. To assert that we can know the denials of skeptical hypotheses also seems bad, reflecting an

unwarranted confidence in our epistemic abilities—these hypotheses are crafted, after all, so that we cannot acquire evidence that would rule them out. Yet to concede that we do not even know that we have hands appears to capitulate to radical skepticism.

Historically, hinge epistemology has been motivated by the traditional epistemological project of addressing skeptical problems. However, it does so in a highly unusual way, by rejecting one of the core assumptions of traditional epistemology—namely, the Universality Thesis, and the associated evidentialist idea that the legitimacy of all believing is wholly determined by one’s evidence. Hinge epistemology comprises a family of views that seek to develop a systematic epistemic theory from Ludwig Wittgenstein’s remarks in his finally notebooks, posthumously published as *On Certainty* (1969), especially the series of passages from which the approach gets its name:

That is to say, the *questions* that we raise and our *doubts* depend on the fact that some propositions are exempt from doubt, are as it were like hinges on which those turn. That is to say, it belongs to the logic of our scientific investigations that certain things are *in deed* not doubted. But it isn’t that the situation is like this: We just *can’t* investigate everything, and for that reason we are forced to rest content with assumption. If I want the door to turn, the hinges must stay put (*On Certainty* §§341-343)

The core idea is that all epistemic evaluation is possible only relative to certain ‘hinge commitments’ that must be held fixed, beyond doubt. But unlike foundationalist and neo-Moorean responses to skeptical problems, these commitments are not presented as supported by evidence or as being somehow self-evident. Rather, hinges themselves incapable of rational epistemic support, *precisely because* they are held with an optimal degree of certainty. At the same time, the denial that hinges lack rational epistemic support is not intended as a capitulation to the skeptic—these commitments are supposed to support our ordinary epistemic practices. While the various hinge epistemologies have been developed from Wittgenstein’s remarks disagree sharply in their details, they are all committed to at least the following two ideas (Ashton 2019):

- A) *Lack of Evidential Support*: There are some basic intellectual commitments that we take to comprise knowledge, for which we lack evidential support (in some sense).
- B) *Legitimacy*: Nevertheless, it is epistemically legitimate (in some sense) for us to hold these commitments.

As I prefer to think of it, hinge epistemology draws out some anti-skeptical consequences from a general principle (which we find in *On Certainty*) of how some commitments provide rational support for others:

*The Rational Grounding Principle*: In order for *S* to have rational grounds to doubt (or believe) that *p*, those grounds must themselves be initially more subjectively certain for *S* than the proposition *p* to be doubted (or believed) is for *S*.<sup>1</sup>

This principle captures some of the intuitive force behind the closure and transmission principles. In order for one’s epistemic warrant for one’s belief that *p* to transmit and so

---

<sup>1</sup> A version of this principle is discussed, seemingly favorably, in Pritchard (2016). However, Pritchard has recently clarified that while he finds this principle in Wittgenstein’s work, he does not endorse it himself (2019 S3658), in response to a comment by Matt Jope (2019). I shall however make use of the Rational Grounding Principle in explicating my preferred approach to hinge epistemology.

support belief that  $q$ , it seems important that the warrant one has for  $p$  is independent of and initially more certain for one than any warrant one already has for  $q$ —otherwise, the warrant that is transmitted would be redundant in providing epistemic support for  $q$ . In addition, if the support for one’s knowledge that  $p$  is to allow one to extend one’s knowledge through a competent deduction to  $q$ , one must initially be more certain of  $p$  than of  $q$ —or else no *new* knowledge will be acquired.

With the *Rational Grounding Principle* in hand, we can see why hinge commitments—as propositional attitudes characterized by an optimal degree of certainty in the truth of some proposition—are simultaneously rationally groundless and also incapable of directly epistemically supporting other beliefs. Because hinge commitments are already optimally certain within an epistemic framework, there can be no *more certain* commitments that could serve as their rational grounds—hence, hinges lack evidential support. (Nor could there be any more certain commitments, incompatible with the hinges, that could be used to call them into doubt). Yet, being themselves rationally groundless, hinges seem to have no warrant that could be transmitted to other beliefs via transmission principles. This leads to the conclusion, shared by many hinge epistemologists, that hinges are in some way removed from epistemic evaluation.

Rather than standing directly in relations of epistemic support to other propositions, Wittgenstein portrays hinges as instead having the status of a *rule* or a *standard*: They play the role in an epistemic system of determining what can count as evidence for what. They provide “the inherited background against which I distinguish between true and false” (*OC* §94), and determining what is to be tested by what (*OC* §125). As such, hinges themselves are removed from direct epistemic evaluation. It is thus inappropriate for them to appear in instances of the closure principle as it is exploited by the skeptical problem. The attempt to ‘extend’ one’s knowledge via known entailment, by deducing from the proposition that I have hands, and my knowledge that if I have hands, I am not a handless BIV, the conclusion that I am not a BIV, one illegitimately treats what is in fact a standard for determining what can count as evidence for what as itself something that can be evidentially supported or undermined. This is the key to the hinge epistemic dissolution to skeptical problems; the problem turns on an *unrestricted* application of closure and transmission principles, whereas properly understood, some commitments—such as my commitment to the proposition that I am not a handless brain-in-a-vat—lie outside the scope of such evaluation. It is in this sense that hinge epistemology rejects the Universality Thesis: not all commitments which it is legitimate for one to hold are subject to epistemic evaluation in terms of rational support (Pritchard 2016).

While there has been a consistent and growing interest in hinge epistemology over the past two decades or so, the plausibility of the view remains contested. Hinge epistemology is in some ways still a fringe view in contemporary epistemology, perhaps because it rejects principles central to traditional epistemology. I now turn to consider one of the major challenges that has been raised for hinge epistemology in general—the Demarcation Problem (Ohlhorst 2021, 2022). After discussing the challenge, I will show (in the rest of the paper) how a salience-based approach to epistemology can be naturally paired with hinge epistemology to address it.

In the philosophy of science, the demarcation problem is the problem of distinguishing between genuine science and pseudoscience (Popper 1968). Applied to hinge epistemology, the Demarcation Problem is the problem of how to distinguish between genuine and counterfeit hinge commitments. Hinge commitments, we have seen, are supposed to be legitimate to hold despite lacking evidential support. However, clearly not all commitments lacking evidential support are legitimately held. We need only consider dogmatically held beliefs, delusions, wishful thinking, fixed ideas, and possibly prejudicial attitudes to see that this is so. The challenge for hinge epistemology then is to explain what separates genuine hinges—paradigmatically including propositions such as ‘There is an external world’—from what Ohlhorst describes as ‘bizarre’ and ‘delusional’ commitments that might be held by some individuals to a high degree of certainty, yet are supposed to be clearly false and incapable of grounding good epistemic practices—such as firm commitment to the proposition that the earth is flat, or that one must knock on ladders with a particular rhythm to prevent them from melting (2022, p. 324).

Differentiating hinges from other forms of highly certain belief is especially difficult given that there are no clear restrictions on what sorts of propositions or contents can play the hinge role. Broadly foundationalist approaches in epistemology (going back at least to Descartes’ *cogito*) appeal to properties of certain propositions or kinds of content as making them especially epistemically secure—e.g. logical or mathematical truths, *a priori* truths, perceptual seemings, basic self-knowledge of one’s current mental states. Some hinge epistemologists have likewise restricted their views to countenance only certain kinds of propositions as legitimate candidates for hinges, and might hope to avoid the demarcation problem this way (Coliva 2015; Wright 2004). However, such restrictions on the possible contents of hinge commitments does not cohere well with the startling range of examples Wittgenstein provides—including commitment to the proposition that one’s friend does not have sawdust in his body or head (§281), that if someone’s arm is cut off it will not grow again (§274), that motor cars do not grow out of the earth (§279), that I am now living in Oslo (cf. §421), or that the water in the kettle on the gas flame will not freeze but boil (§621)—these examples seem deliberately heterogenous, eschewing categorization under some single type of especially indubitable content; none of these are standard foundationalist fare. (Note, also, that these sorts of commitments are ordinarily things which we do not attend to or find salient—more on this later). Moreover, the approach to hinge epistemology I’ve sketched here operates without any assumptions about the content of hinge commitment. In principle, one could have a hinge commitment toward any proposition, as long as it is help to an optimal degree of certainty and plays the role of setting a standard of what can count as evidence for what. This seems to exacerbate the demarcation problem. What could possibly distinguish such commitments from ordinary beliefs (such as my belief that it will snow today) in an epistemic framework?

In the next section, I sketch an outline of what a salience-based approach to epistemology would look like and how it differs from the standard belief-based approach in epistemology. Then I use a salience-based approach to give a precise account of what differentiates hinge commitments from other kinds of firm belief in terms of the topography of salience maps. In short, the idea is that hinges are not demarcated by a certain kind of content or intrinsic indubitability, but rather by their function in an epistemic framework, a function that can be modeled in terms of their location in a salience structure. Hinges help to determine the boundaries of what is salient, hence to shape the direction of inquiry, including what topics are and are not apt for investigation. Appreciating the place of hinges in salience structures will help us see more concretely the role they play in enabling ordinary belief and doubt.

### 3. Saliency Structures

Saliency itself has only very recently become a topic of direct philosophical concern. To say that something is salient to you is to say that it ‘stands out’, that it calls for your attention. Minimally, then, we can understand saliency as the propensity or tendency for something to draw attention. Corresponding to this, a saliency *structure* can be understood as a way of ranking or sorting items according to the degree to which they are apt to draw attention. Drawing on recent work in the normative and epistemic dimensions of saliency, I will argue that saliency structures provide a useful tool for modeling the role of hinge commitments in epistemic evaluation.

There are some important distinctions to keep in mind as we explore the relevance of saliency to epistemology. First: we can distinguish between objective and subjective senses of saliency: to say that something is objectively salient is to say that it makes a demand on your attention, regardless of whether you do in fact attend or are disposed to attend to it. In the subjective sense, saliency is just your tendency or disposition to attend (Archer 2022, 114). Second: One can distinguish between external (object-centered) and internal (or psychological) approaches to saliency; from an external approach, it makes sense to ascribe saliency as a property of non-mental objects or events (in addition to mental ones). For instance, we can say that those neon green boots themselves are salient—their striking color makes them apt to draw attention in a shoe store of mostly grey and black boots. (Munton 2021 favors the externalist approach to saliency). From an internalist approach, what is salient are not properly objects in the external environment, but rather psychological elements—strictly speaking, on this approach, it is not the *boots* (or their color) that is salient, but your experience of them that is salient. Your impression of the bright green color stands out against the other more mundane aspects of your current experience and mental goings-on. (Watzl 2017 favors the internalist approach). In this paper, I will use Munton’s approach to saliency structures (2021) as a framework for thinking about hinge commitments.

Munton presents a view that is externalist, in that it encompasses both mental and non-mental items as possible elements in saliency structures. It also incorporates both subjective and objective approaches, acknowledging that both the interests, concerns, and habits of the individual, as well the design of the external environment and features of what is happening at a given moment matter to determining what is salient. More precisely, in Munton’s view, a saliency structure is “a function from pairs of tasks and contexts, to a set of accessibility relations, which describe the subjects dispositions to attend in that context and with that task, or set of tasks” (ms., 29). For convenience, however, saliency structures can just be thought of as sets of accessibility relations constituted by the subjects dispositions to access information.

Let us unpack this account a bit. Items in the subject’s mental and physical environments encode information. Attending to those items is a way of accessing the encoded information. The accessibility of a piece of information is a disposition—which in Munton’s framework is just a tendency or probability—of the subject to attend to an object encoding that information. Thus, information encoded by objects that a subject is highly likely to attend to is more accessible than information encoded by objects that a subject is less likely to attend to. How likely a subject is to attend to a given object can only be determined relative to the subjects tasks and context—this is why Munton prefers to speak of *centered* networks of accessibility relations. Depending on your context (for instance, whether one is in a library or hiking through a forest) some information will be more easily accessed relative to other information; additionally, depending on your task (e.g. researching saliency structures, or finding a path home) some information will be more rewarding to access in one’s context



relative to other information. Both factors (context and task) together determine the accessibility (probability of accessing) information. Accessing information about what others have to say concerning salience structures is highly unlikely if I am in the woods trying to find the path home; however, it is much more likely that I will access such information if I am in a university library and trying to research salience structures. These two factors—context and task—seem to correlate, respectively, to (i) the underlying arrangement of information in an environment (an arrangement determined by a context), and (ii) the costs and benefits of accessing certain information in that environment (relative to task). Accessing information about salience structures is rather costly when one is scanning the trees ahead for markings one recognizes, since this information is a long distance away from one's position in a salience structure; but if the pay-off of accessing such information is high enough, in terms of the benefits it brings to completing a task (say, it helps to solve a puzzle in a paper one has been writing), one might nevertheless be disposed to access such information.

A useful analogy Munton uses to explicate the various factors involved in salience structures is through appeal to maps and transportation systems, such as a map of the London Underground network. We can think of such a map as centered—you have a particular location in it (sometimes marked “You are here”). The connections between stops represent accessibility relations between locations; from stop A, location B can be more accessible than location C, while from stop D, location C might be more accessible than location B. This sort of structure is apt, for instance, in describing how information about salience structures is, relative to the location (task-context pair) of being in a library trying to research salience structures, more accessible than information about the best path through the woods to get home from some spot on a hill. By contrast, when starting from a location of being in the woods trying to find a path home, information stored in one's memory about various landmarks, for instance, is much more likely to be accessed than information about salience structures. Accessing information about salience structures while walking through the woods might be more costly than accessing such information while in a library (there is a greater informational space to traverse to get there), but depending the potential payoff with regard to one's tasks, it can also bring a large enough benefit to offset these costs.

Salience, then, is influenced by multiple factors. One important aspect is the underlying informational landscape. This is partly an external factor, insofar as one's external environments make some information more easily accessible than others (libraries contain different sorts of information than forests). But it is also partly internal, shaped by the interests, concerns, and habits of the individual (Munton ms., 26-27). Learned associations between concepts, for instance, makes it so that when one concept is activated, associated concepts are also more easily accessible (see comparison between salience structures and semantic memory, Munton, ms., 41-43). Internal aspects of salience structures can also change over time: “the potential reward of accessing information over time reshapes the underlying accessibility of that information” (Munton, ms., 29). As one comes to be seriously interested in chess, for instance, one will become much more likely to attend to chess-relevant information in their environment (e.g., they will notice if there is a chessboard in the corner of a room, whereas a non-chess-obsessed person would be less likely to notice this). Another other important factor influencing salience is the current task of the individual. The factors shaping internal information landscapes are closely related to the other major factor influencing salience structures—the tasks, concerns, goals of an individual at a time. As inattentional blindness experiments reveal, the tasks one has shape what information in the environment one is disposed to attend to, and what information gets ignored.

We are now in a position to begin describing the significance of salience structures for thinking about hinge commitments. An initial caveat that needs to be made is that, whereas most working in hinge epistemology assume a propositionalist framework for thinking about hinges (Moyal-Sharrock 2004 is an important exception), salience structures are plausibly not best thought of as propositionally structured, or as representational. An ordering of information by accessibility does not represent anything as being the case, though the information so ordered can itself sometimes take a representational guise. But I do not mean to claim that hinges *are* salience structures. Rather, the claim is that hinges can be elements of salience structures; they figure among the things one could attend to, that can be sorted in terms of relative accessibility given a task and a context.

What I wish to claim is that hinges, by their nature, occupy a special place in salience structures. In the previous section, I proposed a view on which hinges are (i) heterogeneous with respect to their content; (ii) play the role of a rule or standard by which one determines what is to count as evidence for what, where; (iii) they can play this role because they are held to an optimal degree of certainty, and; (iv) as a result, they are themselves immune to direct epistemic evaluation in terms of evidence. It is significant that examples of hinges are typically sentences of which we are so certain, we ordinarily do not explicitly entertain them in thought and conversation; they provide the assumed background against which ordinary epistemic practices take place. That is: hinges are characterized by an exceptionally *low* degree of salience across a broad range of contexts and tasks. They are almost never salient; they only stand out to one in very unusual circumstances. A paradigm example of hinge commitments (discussed in Wittgenstein and onwards) is my commitment to the proposition that I have two hands. (*On Certainty* opens with a comment on the relation of this hinge certainty to the rest of our ordinary epistemic practices: “If you do know that here is one hand, we’ll grant you all the rest” §1). While we might regularly have some reason to attend to our hands, it is only very rarely salient to entertain the thought that one has hands—in practice, it might only ever come up when discussing radical skepticism.

There is a reason that hinges will tend to have a low degree of salience. Attention and salience are often viewed as involving the ordering of information. Information, according to the influential Shannon Information Theory, reduces uncertainty: “when a signal Y provides us with information about some random variable X, it does so by reducing our *uncertainty* or *surprise* with respect to X” (Wu 2014, 275). Some information can reduce uncertainty (with respect to an epistemic goal) to a greater extent than other information, and thus counts as more relevant (to an epistemic goal). Given this approach to information, it is epistemically optimal for salience structures to order information such that information that is expected to reduce uncertainty to the greatest extent is prioritized over information that is expected to reduce uncertainty to a lesser extent. Because hinge commitments reflect what one already takes to be certain to an optimal degree, any information that bears directly on a hinge commitment can be expected to be irrelevant—there is no uncertainty to be reduced. That which one is most certain of, one has least reason to directly investigate.

Not everything that has a low degree of salience across a broad range of contexts and tasks constitutes a hinge, however. Hinges have a low degree of salience due to the lack of uncertainty to be reduced. Other objects can have a low degree of salience not because we are especially certain about them, but because investigating them is usually not relevant to epistemic projects that matter to us. The dead AA battery at the bottom of the junk drawer in my apartment is almost never salient for me. The battery itself, as possible object of attention, is not propositionally structured, hence not truth-apt, hence not something about which I can have a firm alethic commitment. Yet even when considering propositions concerning that

battery (such as “There is a dead AA battery at the bottom of my junk drawer”), such propositions are not necessarily associated with any firm alethic attitude; I might not be at all certain that there still is or ever was such a battery there. The presence of a dead AA battery in my junk draw will tend to have a low degree of salience, intuitively, just because I am very unlikely to have any epistemic projects for which information about what is at the bottom of my junk drawer will be relevant. There is uncertainty to be reduced here; it is just not an uncertainty that matters for me, practically or epistemically, in most contexts.

Since hinges are, by their nature, nearly always irrelevant to one’s epistemic tasks, salience structures do not have to be designed to facilitate access to them (this is another way to say that they are rarely if ever prioritized). We can illuminate the special status of hinges again by appeal to maps—this time a topographic map, not a transportation system. Topographic features of a landscape naturally shape how accessible some locations in the area are from other locations, and so tend to guide how buildings, neighborhoods, and transportation systems in the area are designed. In such a map, hinges can be thought of as the peaks of mountains or steep hills; they shape how cities and towns are built around them, but typically are not themselves suitable as a central hub of a transportation system.

In the context of investigating the role of hinges in salience structures, it is important to distinguish how a piece of information can be *relevant* (or not) from it’s being *important*, or *matter*ing to the individual. It is easy to run these together, and to think that insofar as something is relevant, it is important or significant. Often these will together, but being relevant and being a significant concern for an agent can come apart. This is because, first, something might be relevant though insignificant, because one can have cognitive projects that do not deeply reflect what one cares about as a person. And second, something can be highly significant in one’s epistemic framework and one’s life goals in general, even though it is never or only rarely relevant for one’s cognitive projects. The second point is especially relevant in connection with hinge commitments.

Hinge commitments often reflect things that matter deeply to individuals, yet the hinges themselves are not usually relevant to the cognitive tasks that come and go. Sophie Archer makes draws out this point explicitly, acknowledging that:

even if we were to imagine a codification on which you evaluate your partner as mattering most, followed by your work etc. it is clear enough that this will not always correspond with the degree of salience each has for you at any given time. Your partner is not *always* the most salient aspect of your experience and/or thought. In fact, precisely in virtue of their ‘hinge-like’ status in your life, your pressing concerns, such as the work you need to be getting on with that day, will often usurp them in this role, as will someone much less significant to you, generally speaking, who happens to be speaking at that moment (Archer 2022, 124)

This line of thought suggests a close relationship between hinge commitments and *convictions*, where a convictions can be understood as “commitments that have become so woven into our self-narrative that they become part of our self-identity” (Lynch 2019, 59). Convictions, like hinge commitments, enjoy a high degree of certainty, and to frame what lines of inquiry are deemed relevant. Someone with a firm religious conviction will be more likely than someone lacking it to appeal to religious sentiments in making sense of life events, for instance, in a tendency to see providence where others see coincidence, etc. However, convictions are also usually very salient to those who have them across a relatively broad arrange of tasks and contexts. The main contrast between convictions and hinge commitments is primarily in the salience of the commitment to the individual; whereas

convictions might be characterized as *highly* salient across a wide variety of tasks and contexts, hinges are *anti*-salient across a wide variety of tasks and contexts, forming a backdrop of unquestioned certainty.

But over time, commitments that represent major life projects or reflect a major aspect of one's identity can paradoxically shift from having a central place in one's salience structure to take on more the role of a hinge commitment, in the way Archer describes. Attention and salience are inherently dynamic activities; one's attention tends to be grabbed by what is novel, what stands out against a relatively constant backdrop. Sustaining attention on a single task for long periods often requires effort. To make a single idea or object the animating concern of one's mental life for an extended period of one's life is rather difficult. What starts as a conviction, something infusing one's epistemic projects in general and taking on a high degree of salience can, through habituation, perhaps transition into a hinge commitment. The conjecture here is that inculcating convictions is one route to acquiring a hinge commitment.

Hinges are unusual in that they enjoy both a low degree of salience and an optimally high degree of certainty, providing a background terrain around which salience structures are oriented. This allows us to distinguish hinges from highly certain beliefs that are also highly salient under certain task-context pairs. For example, consider the following *Sports Fan* case:

*Sports Fan:* Jill is a sports fan, an avid follower of the Chicago Bulls. Jill proclaims a high degree of certainty in the proposition that the Bulls will win their upcoming match against the Boston Celtics. The forthcoming victory is highly salient to Jill in the week leading up to the match. Any conversation with Jill is likely to lead to basketball-related topics, including statistics about scoring, rebounds, etc., of players on the two teams, which Jill cites as evidence of the Bulls' certain victory.

While Jill claims that the Bulls' victory in the upcoming match is utterly certain, this commitment does not have the characteristics of a hinge commitment. The prioritization of this commitment in Jill's salience structure reflects the relevance of Bulls-related information for the cognitive tasks Jill undertakes, whereas hinge commitments tend to take a background role in shaping salience structures, such that they generally go without notice, at least until philosophers start raising skeptical hypotheses.

Characterizing hinges as having a high degree of certainty plus low salience across a broad range of task-context pairs thus goes some way toward addressing the demarcation problem, since it allows us to distinguish hinges from some other commitments that might initially seem hinge-like (e.g. convictions). To fully address the demarcation problem, though, we need not just to give suitable psychological criteria, but to explain how these features relate to the distinctive epistemic role of hinges. What is the epistemic significance of the fact that hinges have very low salience? I turn to this question in the next section.

#### 4. The Epistemic Role of Hinges in Salience Structures

As Munton points out, salience structures do not figure cleanly fit into the framework of traditional epistemology, making it difficult initially to see how they might be epistemically evaluated. Nevertheless, it is clear that salience structures have epistemic significance, since they shape the accessibility of information, and hence play a role in the collection of evidence and belief-formation.

In general, salience structures are epistemically good, according to Munton, when they are *flexibly relevant*. That is, a good salience structure is such that, for whatever particular cognitive task you are engaged with, the salience structure makes relevant information for

that task salient. The flexibility resides in the point that different information is relevant to different tasks, and so a good salience structure facilitates transitions in the salience of information such that no matter what task one has one is likely to access the most relevant information. As much as possible, then, a salience structure should be like a transportation system that makes it easy to get from any given starting point to any other location in the system.

An epistemically well-ordered salience structure enables the subject to be sensitive to evidence relevant for her epistemic projects; given that evidence  $E$  has direct bearing on inquiry  $I$ , then, were  $I$  to be an inquiry the subject undertakes,  $E$  comprises information relevant for the subject, which her salience structure should then prioritize. With this account of the epistemic evaluability of salience structures in hand, we are better able to understand how it is that hinge commitments work in shaping evidence.

Hinge commitments are supposed to play a particular role in epistemic evaluation; they set a standard or rule for determining what counts as evidence for what, a role that they can play because they are themselves held to an optimal degree of certainty (Coliva 2015; Pritchard 2016). As Wittgenstein puts it: “My having two hands is, in normal circumstances, as certain as anything I could produce in evidence for it. This is why I am not in a position to take the sight of my hand as evidence for it” (*OC* §250); and “If a blind man were to ask me ‘Have you got two hands?’ I should not make sure by looking. If I were to have any doubt of it, then I don’t know why I should trust my eyes. For why shouldn’t I test my eyes by looking to find out whether I see two hands? What is to be tested by what?” (*OC* §125). Passages such as this are suggestive of the Rational Grounding Principle discussed in section 2.

Framed in terms of salience structures, we can put the point as follows: hinges, in virtue of their optimal certainty, are (i) not directly relevant to epistemic projects (they don’t help to reduce uncertainty any further), yet (ii) indirectly shape investigation by placing certain branching paths of inquiry out of bounds—namely, inquiry into those propositions incompatible with the hinge commitment. In effect, hinges generate a permissible form of negative epistemic space, ruling out certain inquiries, and hence helping to guide inquiry into epistemically more promising avenues. For instance, by having the proposition that an external world exists as a hinge commitment, one is entitled to take perceptual experiences as providing evidence about happenings in the environment, without first needing to rule out the possibility that those experiences are (say) merely the products of an evil demon. That the external world exists, as a hinge commitment, is not directly relevant in epistemic evaluation, and hence skeptical hypotheses calling it into doubt can be epistemically appropriately disregarded as one seeks to learn about one’s surroundings via perception. In short, skeptical hypotheses such as the BIV or evil demon hypothesis are never relevant to perceptually-guided inquiry.

Interpreted in terms of salience structures, hinge epistemology can be seen to provide a version of the relevant-alternatives approach to knowledge, where radical skeptical hypotheses are basically never relevant alternatives. From this perspective, the skeptic, by raising BIV-hypotheses, evil demon hypotheses, and the like, attempts to make these hypotheses relevant, and by doing so introducing the demand that we be able to rule these hypotheses out in order to be warranted in taking ordinary perceptual experience as evidence about happenings in an external world. The important twist that the hinge epistemic approach adds is that, *contra* David Lewis’s *Rule of Attention* in relevant alternatives theory (1996, 559), merely attending to some alternative is not sufficient to make it epistemically relevant. By raising skeptical hypotheses, we epistemologists do indeed bring attention to alternatives that are incompatible with the ordinary knowledge we usually take ourselves to have, and for

which we cannot acquire conclusive counter evidence to rule out. Merely raising such a hypothesis does not make it relevant, however, because such hypotheses target the very commitments that set a standard for determining the relevance of information to epistemic projects (in this sense, these ‘doubts’ are self-defeating). Radical skeptical hypotheses are properly ignored.<sup>2</sup>

One of the advantages of a salience-based epistemology in general is that salience structures allow us to model not only the rationality of belief, but also how we epistemically evaluate *gaps* in belief. That is, salience structures, unlike traditional epistemological tools, allow us to directly evaluate ‘negative epistemic space’, or *ignorance* (Munton, ms., chapter 4). When I am ignorant of a piece of information, this is because I have not attended to (the object that encodes) it. Not all cases of ignorance are epistemically blameworthy or bad. In some instances, ignorance about whether *p* (in the form of a failure to form a belief about *p*) is permissible, even when one has evidence speaking for or against *p*. One reason for this is just that we have so much evidence available to us, often on matters that are unrelated and so not relevant for our epistemic concerns. For example, I have perceptual evidence available to me now that would justify me in forming a belief about how many trees are visible from my window. Yet I have not formed such a belief, and nor am I epistemically blameworthy for failing to do so, because this information is not relevant to my epistemic projects. What hinge epistemology reveals is that there is another source for proper ignorance: possibilities are properly epistemically ignored when making them salient would undermine the conditions for epistemic evaluation.

In her salience-based approach to ignorance, Munton proposes that ignorance is epistemically bad when it is *resilient*: That is, when one would continue to remain ignorant about *p* even if information about *p* were to become relevant. This principle regarding the permissibility of ignorance falls out of Munton’s broader view about how to epistemically evaluate salience structures. Given that a salience structure is epistemically good insofar as it organizes information in a way that is flexibly relevant—i.e., such that information relevant to current epistemic tasks-in-context are prioritized—a salience structure that fails to prioritize information related to *p*, were information about *p* to become relevant, fails to meet the flexible relevance norm.

One might worry that hinge epistemology supports the rationality of resilient ignorance about radical skeptical hypotheses—possibilities inconsistent with one’s hinge commitments are properly ignored across a wide range of tasks and contexts. Strictly speaking, however, hinges do not really vindicate resilient ignorance, for the possibilities they exclude are really almost never relevant, and so it is not an epistemic fault of salience structures if they fail to prioritize information about these skeptical possibilities. At any rate, it is worth noting that while hinges are normally not going to be salient, this cannot directly generate a worry that hinge epistemology endorses resilient ignorance. To have a hinge commitment to the proposition that *p* is to have an attitude towards *p* that is not agnostic with respect to *p*’s truth (or, for that matter, with respect to skeptical alternatives to *p*). One does not suspend judgment when one has a hinge commitment. Thus, even if one’s hinge commitments are not salient to one, they also do not constitute ignorance in the sense of failing to have an opinion about the relevant matter.

## 5. Conclusion: A Limitation on Norms for Salience Structures

---

<sup>2</sup> For an alternative take, see Gardiner (2020), who advocates for a relevant alternatives approach that takes the rule of attention seriously, and employs it in an explanation of certain forms of epistemic injustice.

In this paper, I have considered how a salience-based approach to epistemology can help make sense of the role of hinge commitments in epistemic evaluation. In brief, the main points are that:

- i) Hinge commitments, in virtue of being optimally certain, are almost never relevant to inquiry—directly epistemically evaluating them is unlikely to reduce uncertainty.
- ii) Hinge commitments, by being anti-salient (or ‘inconspicuous’), set boundaries on inquiry. Hinges define a ‘negative epistemic space’ of lines of inquiry that are ignored—including inquiries about skeptical hypotheses incompatible with particular hinge commitments.
- iii) Although hinges are almost never relevant to inquiry, they play a central role in determining what lines of inquiry will be considered ‘live’ and relevant to subjects. As such, hinges will tend to track what is significant to individuals, despite not themselves being the target of sustained inquiry.
- iv) Characterizing hinge commitments in terms of their place in salience structures also helps to identify features of such commitments that can be used to differentiate them from other sorts of firm commitment, thereby contributing a response to the ‘demarcation problem’ for hinge epistemology.

The focus of this paper has been to identify some ways that adopting a salience-based approach to epistemology can be useful for illuminating the notion of a hinge commitment. In concluding, I would like to briefly consider a connection running in the other direction: implications of hinge epistemology for salience-based epistemology. In particular, I suggest that taking hinge epistemology seriously requires acknowledging an important limitation in how salience structures can be epistemically evaluated.

One of the main conclusions of hinge epistemology is that epistemic evaluation is necessarily a local affair. All epistemic evaluation takes place relative to hinge commitments that are themselves immune to both rational support and to doubt. A ‘globalist’ about epistemic evaluation, by contrast, thinks that any commitment is in principle evaluable for its epistemic credentials. Evidentialists, for instance, are globalists of this kind, insofar as they hold, with Clifford, that “It is wrong always, everywhere, and for anyone, to believe upon insufficient evidence” (1877). Skeptical arguments exploit such global epistemic principles, by pointing out that for basically any belief one would ordinarily take to comprise knowledge, we can identify some entailment of that belief that one’s evidence is in principle incapable of ruling out. Hinge epistemology dissolves skeptical problems arising from such arguments by identifying a kind of commitment that is insulated against direct epistemic evaluation (either positive or negative) and supports ordinary epistemic practices. Epistemic evaluation occurs locally, within the practices so supported.

The limitation to local epistemic evaluation also applies to epistemic norms on salience structures. Munton is clear, of course, that epistemic norms on salience structures are non-evidential; it is possible to epistemically assess salience structures in ways that go beyond mere sensitivity to evidence, and it is sometimes epistemically permissible to fail to form beliefs even when one has evidence that support those beliefs. Munton proposes instead that salience structures should be evaluated with respect to the *flexible relevance* norm. A salience structure is epistemically normatively good insofar as it facilitates access to maximally task-relevant information across a maximally wide range of tasks (ms.). As a regulative ideal for salience structures, the ultimate epistemic goal to aim for (but that we mortals will never attain), Munton describes ‘the Big Knower’: an ideal epistemic agent who has completed the

Big Inquiry, and so has completed all sub-inquiries, thereby coming to know the nature of the universe in all its detail.

The Big Knower, as a regulative ideal, is of course not an achievable goal for any of us among the epistemic ‘fallen’. However, if we take hinge epistemology seriously as an account of the structure of epistemic evaluation, I think we are left with some doubts about the Big Knower as comprising a coherent ideal by which to steer our ship. Given that all epistemic evaluation is in principle ‘local’, such that some commitments must lie outside the scope of rational evaluation (positive or negative), the idea of an ideal knower, one who has well-founded beliefs in all true propositions, and well-founded doubts about all false propositions, becomes incoherent. Such an epistemic position would require a ‘view from nowhere’, and as such, would fail to comprise an epistemic viewpoint at all, rather than standing as the ideal limit for us to strive for.

## References

- Archer, S. (2022). “Salience and What Matters”. In S. Archer (ed.), *Salience: A Philosophical Inquiry* (pp. 113-129). New York, NY: Routledge.
- Ashton, N.A. (2019). “The Case for a Feminist Hinge Epistemology”. *Wittgenstein-Studien*, 10(1): 153-163.
- Clifford, W.K. (1877). “The Ethics of Belief”. In *The Ethics of Belief and Other Essays*. Amherst, NY: Prometheus Books.
- Coliva, A. (2015). *Extended Rationality: A Hinge Epistemology*. Palgrave Macmillan.
- Conee, E.B., and Feldman, R. (2004). *Evidentialism: Essays in Epistemology*. Oxford: Oxford University Press.
- Fricker, M. (2007). *Epistemic Injustice: Power and the Ethics of Knowing*. Oxford University Press.
- Gardiner, G. (2020). “Relevance and Risk: How the Relevant Alternatives Framework Models the Epistemology of Risk”. *Synthese*, 199(1-2): 481-511. <https://doi.org/10.1007/s11229-020-02668-2>
- Jope, M. (2019). “Closure, Credence, and Rationality: A Problem for Non-Belief Hinge Epistemology”. *Synthese*, 198(15): S3565-3575. <https://doi.org/10.1007/s11229-019-02153-5>
- Lewis, D. (1996). “Elusive Knowledge”. *Australasian Journal of Philosophy*, 74(4): 549-567. <https://doi.org/10.1080/00048409612347521>
- Lynch, M. (2019). *Know-It-All Society: Truth and Arrogance in Political Culture*. New York, NY: WW Norton.
- Moyal-Sharrock, D. (2004). *Understanding Wittgenstein’s On Certainty*. New York, NY: Palgrave-Macmillan.
- Munton, J. (2021). “Prejudice as the Misattribution of Salience”. *Analytic Philosophy*, 64(1): 1-19. <https://doi.org/10.1111/phib.12250>
- Munton, J. (ms.). *Priority and Prejudice: The Epistemology of Salience and Attention*.



- Ohlhorst, J. (2021). "Epistemic Austerity: Limits to Entitlement". *Synthese*, 199(5-6): 13771-13787.
- Ohlhorst, J. (2022). "Is There a Problem of Demarcation for Hinges?" *International Journal for the Study of Skepticism*, 12(4): 317-330.
- Popper, K. (1968). "Conjectures and Refutations". *British Journal for the Philosophy of Science*, 19(2): 159-168.
- Pritchard, D. (2016). *Epistemic Angst: Radical Skepticism and the Groundlessness of Our Believing*. Princeton, NJ: Princeton University Press.
- Pritchard, D. (2019). "Scepticism and Epistemic Angst, Redux". *Synthese*, 198(15): S3635-S3664. <https://doi.org/10.1007/s11229-019-02504-2>
- Steup, M., and Neta, R. (2020). "Epistemology". In N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy* (Fall 2020 Edition). <https://plato.stanford.edu/archives/fall2020/entries/epistemology/>
- Wittgenstein, L. (1969). *On Certainty*. G.E.M. Anscombe and G.H. von Wright (eds.). D. Paul and G.E.M. Anscombe (trans.). New York: Harper and Rowe.
- Watzl, S. (2017). *Structuring Mind: The Nature of Attention and How It Shapes Consciousness*. Oxford, UK: Oxford University Press.
- Wright, C. (2004). "Warrant For Nothing (and Foundations for Free)?" *Aristotelian Society Supplementary Volume*, 78(1): 167-212.
- Wu, W. (2014). *Attention*. New York, NY: Routledge.
- Wu, W. (forthcoming). "On Attention and Norms: An Opinionated Review of Recent Work". *Analysis*. <https://doi.org/10.1093/analys/anad056>