



New Directions for Awareness Research

Paul Silva Jr.
University of Cologne

Table of Contents

1. Project Overview	2
2. State-of-the-art	2
3. Research objectives	8
4. Methodology	18
References	19

Dear Reader,

The set of research questions outlined below is too large for one person to make timely progress on. I'm more interested in seeing these issues addressed than I am in seeing them addressed *by me*. (Though I do anticipate working on some of them as I have time. See, for example, '[Awareness by Degree](#)'.) So I'm posting this projected research in the hope of inspiring you to explore, in your own way, the networked set of questions outlined below. Should you begin working on these or related issues, feel free to reach out to me at [<psilvajruoc@gmail.com>](mailto:psilvajruoc@gmail.com). As time permits, I'm happy to weigh-in on and support your research projects.

Please forgive typos and errata in what follows.

All the best,

Paul



1. Project Overview

In *The Conscious Mind*, David Chalmers distinguished questions about **awareness** from questions about consciousness, where the study of awareness is the study of how cognitive systems function in relation to the world. In Chalmers' view, the study of awareness falls squarely within the purview of both cognitive science and philosophy.¹ This project aims to achieve two general goals. First, to provide a philosophical analysis that advances our understanding of factual awareness and other associated world-implicating psychological states. Second, apply the results of this advancement to a range of issues in applied epistemology. In particular, this project will explore:

- (i) the nature, structure, and significance of factual awareness as well as other mental states that factive terms in natural languages refer to,
- (ii) the nature and structure of know-how and its relation to awareness,
- (iii) norms for belief, assertion, action, and deliberation, and their relation to awareness, and
- (iv) how an improved understanding of topics (i)-(iii) impact our ability to address problems in: (a) internet epistemology, (b) AI epistemology, (c) collective epistemology, (d) political epistemology, and (e) the epistemology of encroachment.

2. State-of-the-art

The ability of humans to **represent** the world is central to understanding how we survive in the world. For instance, it is your mind's ability to represent *an approaching car* that enables you to avoid colliding with it. In doing this, our minds represent the world in part by representing *particulars and their properties*. But our minds also represent the world by **representing facts** about the world: our minds can represent not only *the car* (a particular) and *its color* (a property of a particular), our minds can also represent *the fact that there is a car with such-n-such color*. And it is primarily (and perhaps exclusively) with facts that we reason and theorize with in deductive and non-deductive ways that allow us to arrive at scientifically well-founded views about the nature of ourselves and our environments. In this way the representation of facts is central to understanding some of humanity's greatest intellectual achievements.²

¹ Chalmers (1996: 29) writes: "Awareness, like other psychological properties, poses few metaphysical problems. ... Certainly, the notion of awareness is not crystal-clear, so there is room for significant philosophical analysis of just what it comes to. Further, there is room for an enormous amount of research in cognitive science, studying how natural and artificial cognitive systems might function in such a way that they are aware. But the outlines of these research programs are reasonably clear. There is little reason to suppose that the normal course of cognitive science, backed by appropriate philosophical analysis, should not eventually succeed."

² There are at least two ways of understanding 'facts' that should be distinguished for present purposes. Sometimes 'facts' are taken to refer to some *particular*, *x*, *exemplifying some property* F. For example, *the computer's being large* and *this sentence occupying space* are facts in this exemplification sense of 'fact'. But this exemplification sense of 'fact' is to be put aside in what follows. In what follows I'm concerned with the sense of 'fact' that involves *that-clauses*, e.g. being aware of (or ignorant of) *the fact that the computer is large*. The property-exemplification sense of 'fact' is sometimes connected to truth-makers, while the that-clause sense of 'fact' is often connected to true propositions—though it need not be (see Hyman 2017). That-clause facts are also more expansive: they can involve facts involving quantifiers, negatives, disjunctions, and the like. In contrast, 'facts' as exemplifications of properties seem more limited to particulars and the properties they instantiate and standing in stative relations to them is conceptually undemanding. In contrast, as pointed out by Dretske (1993) and others, we gain access to (e.g. become aware of) that-clause facts *in part* by deploying concepts that allow us to host thoughts with propositional content. See Mulligan and Correia (2021) for a discussion of senses of 'fact'.



The significance of our ability to represent facts is born out in natural languages, which contain a rich and varied vocabulary of factive terms. On the one hand there are **factive stative verbs**, these are verbs that signal that one has come to be in a mental state that represents a fact (rather than a falsehood) as obtaining. Core examples of factive stative verbs include: ‘knows that’, ‘realizes that’, ‘recognizes that’, ‘notices that’, ‘discerns that’, ‘understands that’, ‘discovered that’, ‘learned that’, ‘perceives that’, ‘sees that’, ‘hears that’, and ‘remembers that’.³ These are called *factives* because they refer to states that require truth. You cannot, for example, know, or discover, or see *that Mont Blanc is large* if Mont Blanc is not large.

In addition to factive stative verbs, natural languages contain **factive non-stative verbs**. These verbs signal that one is *not* in some state that represents a fact as obtaining. Factive non-stative verbs, and verb phrases, include: ‘forgot that’, ‘overlooked the fact that’, and ‘missed the fact that’. These are factives because one cannot, for example, overlook or miss the fact *that this sentence is written in Spanish* because it is not written in Spanish. For this reason, factive non-stative verbs are not mere denials of their stative counterparts.⁴

It is unsurprising that natural languages have such factive verbs. States of knowledge, understanding, and so forth are of enormous practical and explanatory importance. Knowing facts, understanding facts, perceiving facts, etc. helps us achieve our practical and theoretical goals. Further, our failures to achieve our goals are often to be explained in terms of facts that we overlooked or forgot.

That is all reasonably familiar territory. Less familiar is the range of **factive adjectives** and factive adjectival expressions. Peter Unger (1975: 171-176) drew attention to a wide range of *emotive* factive adjectives, all of which are statives.⁵ But there are also many *epistemic* factive adjectives. Here are *some* such expressions:

Statives	Prospective Statives	Non-Statives
<ul style="list-style-type: none"> ● ‘aware that’ / ‘aware of the fact that’ ● ‘conscious that’ / ‘conscious of the fact that’ ● ‘obvious to her that’ ● ‘clear to her that’ 	<ul style="list-style-type: none"> ● ‘obvious that’ ● ‘clear that’ ● ‘it is certain that’ ● ‘evident that’ 	<ul style="list-style-type: none"> ● ‘ignorant that’ / ‘ignorant of the fact that’ ● ‘unaware that’ / ‘unaware of the fact that’ ● ‘oblivious that’ / ‘oblivious to the fact that’

Positive statements involving the **stative** expressions entail that some agent hosts a mental state that represents a fact (rather than a falsehood) as obtaining. For example, it cannot be obvious *to Ralph*

³ While ‘truly believes that’ is an epistemologically central factive expression involving the verb ‘believes’, it is not among the intended class of factive stative verbs as its factivity is owed to its adverbial modifier ‘truly’. None of the other verbs listed above needs a modifier to guarantee factivity.

⁴ Compare ‘it is false that S remembers that p’ with ‘S forgot that p’. If p is false, then the first claim will be true given that ‘remembers that’ is factive. But p’s falsehood ensures that the second claim is false because ‘forgot that p’ is also a factive. So ‘it is false that S remembers that p’ and ‘S forgot that p’ are not logically equivalent.

⁵ His list included: ‘amazed that’, ‘amused that’, ‘angry that’, ‘annoyed that’, ‘ashamed that’, ‘astonished that’, ‘astounded that’, ‘delighted that’, ‘depressed that’, ‘disappointed that’, ‘elated that’, ‘embarrassed that’, ‘excited that’, ‘furious that’, ‘glad that’, ‘grateful that’, ‘happy that’, ‘horrified that’, ‘irritated that’, ‘over-joyed that’, ‘pleased that’, ‘proud that’, ‘sorry that’, ‘surprised that’, ‘thankful that’, ‘tickled that’, ‘unhappy that’, ‘upset that’, ‘impressed that’.



that $2+2=4$ unless: $2+2=4$ and Ralph hosts a representational state that has that fact as content. Similarly, Ralph cannot be aware that water is H_2O unless water is H_2O and he hosts a representational state with that fact as content. In contrast, positive statements involving the **prospective stative** expressions do not entail that any particular agent hosts a representational state with the relevant fact as content.

In contrast, positive statements involving the **non-stative** expressions signal that some agent *fails* to represent a fact as obtaining. Like the non-stative verbs, these factives are not mere denials of their stative counterparts. To be ignorant that, unaware that, or oblivious to p , p must be a fact (Silva & Siscoe 2023).⁶

All of this raises a question: what relation do these factive terms, and the mental states they implicate, bear to each other? Unger (1972, 1975) argued that all the factive stative verbs and the factive stative adjectives are knowledge-entailing, while all the factive non-stative terms entail the absence of knowledge. That is, to stand in any of the factive mental states referred to by these terms requires that one be in a knowledge state; and to stand in any of the relations referred to by the factive non-stative terms requires that one not be in a knowledge state.⁷ Decades later, Williamson (2000) went on to influentially defend a related claim: knowledge is the most general factive stative attitude. On Williamson's view, not only is every factive stative term knowledge-entailing, it is knowledge-entailing because knowledge is the most general factive stative relation. Thus, all the factive stative terms (verbs and adjectives) refer either to mental states that are to be identified with knowledge or are somehow reducible to specific ways of knowing (seeing that p , remembering that p , etc.). Let's use the term **Knowledge-First Reductionism** to refer to the generic idea that these natural language factives are to be *somehow* understood either in terms of knowledge or in terms of the absence of knowledge. Thus, it is part of Knowledge-First Reductionism that some factive relations are to be understood in terms of derivative knowledge-relations like *being in a position to know* or *being capable of knowing*.

Knowledge-first reductionism is of **enormous importance**. The last two decades of research on factive relations has been inundated with hundreds of papers and dozens of books promoting, or else opposing, various elements of 'knowledge-first epistemology'.⁸ At its core, knowledge-first epistemology attempts to enlist one factive stative verb, 'knows', as the central explanatory relation in epistemology: knowledge, it has been argued, can be used to explain the nature of belief, justification, excuses, the possession of reasons, rational action, know-how, norms of assertion, the nature and aims of inquiry, the nature of ignorance, the nature of basic epistemic competencies, and more besides. In addition to that, knowledge-first epistemology has been part of

⁶There is a further distinction to draw that cuts across the factive stative verbs and the factive stative adjectives. Some factives are **amodal** in the sense that they fail to signal *how*, or *the way in which*, one came to be in a given factive state. For example, verbs like 'knows that', 'understands that', 'realize that', and adjectives like 'is aware of the fact that' and 'is clear that' are amodal in that sense. In contrast, there are **modal** factives: 'perceiving that', 'seeing that', 'hearing that', 'remembering that', and 'was informed that'—each of which implicates something about *how* one came to stand in a factive relation to some fact.

⁷ While 'is aware that' was noted as a factive stative adjective by Unger, the class of factive stative adjectives Unger (1975: 171ff) primarily drew attention to were, what we might call, factive stative adjectival **emotives**: 'is amazed that', 'is amused that', 'is angry that', 'is delighted that', 'is elated that', 'is embarrassed that', 'is glad that', 'is grateful that', and 'is happy that'.

⁸ For a pathway into this sprawling literature see the edited volumes of Greenough & Pritchard (2009), Carter, et al. (2017), and Mitova (2018) and the survey articles of Littlejohn (2022) and Silva (2020). As further evidence of the significance of Knowledge-First reductionism, consider that Williamson's *Knowledge and Its Limits* – the work that initiated the deluge of contemporary explorations into knowledge-first epistemology – has more than 6,600 citations on Google Scholar.



many **high-profile grants**. For some examples of recently funded (**and well-deserving!!!**) projects that explicitly involve knowledge-first elements see Prof. Mona Simion's ERC funded project "Knowledge-First Social Epistemology",⁹ Prof. Christoph Kelp and Prof. Anne Meylan's SNF funded project "Evidence: Knowledge and Understanding",¹⁰ and Prof. Maria Lasonen-Aarnio's ERC funded project "Competence and Success in Epistemology and Beyond".¹¹

Resistance to Knowledge-First Reductionism has been explored on the grounds that some factive stative verbs seem to refer to states that have different properties than knowledge. For instance, some have argued that 'sees that' cannot be understood in terms of 'knows that' because seeing-that has a weaker non-accidentality condition than knowledge.¹² If such arguments hold, it's likely that other factive verbs fail to be knowledge entailing as well. After all, if you believe that *p* because you saw that *p* then, arguably, you *learned that p*. So if you can see that *p* without knowing that *p*, then you can also *learn that p* without knowing that *p*. Related opposition to Knowledge-First Reductionism has been explored in connection with *remembering that p* (cf. Bernecker 2010) as well as *understanding that p* (cf. Hannon 2021). We'll return to this below.

So while debates about the relation between knowledge, understanding, seeing, and remembering have been explored to some extent over the past two decades in relation to Knowledge-First Reductionism, what has been largely unexplored is the extent to which Knowledge-First Reductionism is threatened by insights about the mental states that factive adjectives refer to. Indeed, in addition to the arguments involving seeing and understanding just hinted at above, there are **at least five further arguments against Knowledge-First Reductionism** stemming from insights about these adjectives. I'll briefly sketch them.

The first kind of argument against Knowledge-First Reductionism casts doubt on the ability to understand factual ignorance in terms of knowledge (which is the standard view of factual ignorance). To help appreciate the significance of this, consider the very first line of Williamson's *Knowledge and Its Limits*, a foundational text for Knowledge-First Reductionists:

"If I had to summarize this book in two words, they would be: knowledge first. It takes *the simple distinction between **knowledge and ignorance** as a starting point from which to explain other things, not as something itself to be explained*. In that sense the book reverses the direction of explanation predominant in the history of epistemology." (Williamson 2000: v)

In opposition to this, Silva and Siscoe explain in "Ignorance and Awareness" (2023, *Noûs*) that there are many cases where agents fail to know that *p*, but also fail to be ignorant of the fact that *p*. So there seems to be some factive relation *between* knowledge and ignorance. Language gives us a clue as to what that middle ground is. For the factive non-stative adjective 'is ignorant that' can easily be understood as the contradictory of the factive stative adjective 'is aware that'. So the natural pairing to leverage in "reversing the direction" of epistemology is not between knowledge

⁹ Project Website: <https://cordis.europa.eu/project/id/948356>.

¹⁰ Project Website: <https://www.cogito-glasgow.com/evidence-knowledge-understanding>.

¹¹ Project Website: <https://www.helsinki.fi/en/researchgroups/competence-and-success-in-epistemology-and-beyond/project-descripti-on>.

¹² McDowell (2002), Turri (2010), Pritchard (2012b), Schroeder (2015), and Silva (2023).



and ignorance, but between *awareness and ignorance*. Call this argument against Knowledge-First Reductionism, **the argument from ignorance**.¹³

The second style of argument against Knowledge-First Reductionism is explored in detail in *Awareness and the Substructure of Knowledge* (2023, Oxford University Press), henceforth abbreviated as ‘ASK’. ASK focused on the factive stative adjectival expression ‘is aware that’ and carefully explored a wide range of **case-based arguments** that undermine Knowledge-First Reductionism. It demonstrated that there is a wide body of evidence supporting the idea that there are discrete modes of factual awareness (visual, memorial, and self-evident awareness of facts) that can fall short of knowledge and knowledge-centric relations because knowledge is a *more demanding* state than mere awareness of facts. Knowledge requires belief, the absence of objective defeaters, and a robust anti-luck condition. But factual awareness, in contrast, does not require belief, the absence of defeaters, or an *equally* robust anti-luck condition (see ASK Ch.2-4). Many of these case-based arguments that appear in ASK were first defended by other epistemologists. Sven Bernecker (2010) argued that knowing that p is more demanding than remembering that p (which entails memorial awareness of facts); Huemer (2001), McDowell (2002), Pritchard (2012), Turri (2010), and Schroeder (2021) have argued that knowing that p is more demanding than seeing that p (which entails visual awareness of facts); and Huemer (2001) and Silva (2019, 2021, 2023) have argued that inferential knowledge is more demanding than *inferential awareness of facts*.¹⁴ What ASK added to this literature was several novel case-based arguments against Knowledge-First Reductionism, a strengthening of existing case-based arguments against Knowledge-First Reductionism, and new counterarguments to defenses of Knowledge-First Reductionism.

The third kind of argument against Knowledge-First Reductionism explored in ASK is **an argument from generalization**. It is an argument that relies only on very general and widely accepted principles about knowledge. Here’s the sketch. Knowledge is, on all accounts, a kind of *non-accidental true belief*.¹⁵ *True belief* is itself a kind of *true representation*. So there is a more general kind of state of which knowledge is but one instance: *non-accidental true representation*. Since knowledge is but one instance of this more general state, knowledge cannot be identical to it.¹⁶ ASK (Ch. 5) goes on to argue that there is also no way of reducing this more general state to other knowledge-centric relations, e.g. being in a position to know, or being capable of knowing truths. Much of what follows would proceed unhampered if we were simply to identify the awareness of facts with states of non-accidental true representation.

The fourth style of argument against Knowledge-First Reductionism is **an argument from gradability**. Many theorists have observed that knowledge is not gradable.¹⁷ It doesn’t make sense to talk of ‘knowing that p *more than* some else knows that p’ or ‘*somewhat* knowing that p’ or

¹³ However, for the purposes of the paper, Silva and Siscoe took an ecumenical approach that left open the possibility that awareness could be connected to being in a position to know. It is *Awareness and the Substructure of Knowledge* that explores a range of problems with this, and thus closes this gap.

¹⁴ One could, of course, argue that visual and memorial awareness that p is distinct from seeing that p or remembering that p, but it is often assumed that they come to the same thing. For example, Dretske’s (1993) ground-breaking work on factual awareness identifies without comment or argument that seeing that p just is being visually aware of the fact that p.

¹⁵ Beliefs that are *non-accidentally true* are beliefs that are held or formed in such a way that their truth is not a matter of chance in some epistemically relevant sense. A belief’s being reliably formed, or being safely formed, or being sensitively formed, or being justified by the facts one possesses each illustrate ways of being non-accidental in this broad sense.

¹⁶ This argument of ASK also has some precedent in Huemer (2001).

¹⁷ Ryle (1949), Dretske (1981), Stanley (2005), and Brogaard (2016).



‘*fully* knowing that p’. Notably, Hetherington (2001) defended the gradability of ‘knows that’, relying on expressions like ‘knows well’ and ‘knows better’. But Stanley (2005, Ch.2) demonstrated that there is a strong body of linguistic evidence that supports the non-gradability of ‘knows that’, ‘knows well’, and ‘knows better’.¹⁸

In contrast, the factive stative adjectives are gradable. For one example, expressions such as ‘is *somewhat* aware [of the fact] that’, ‘is *vaguely* aware [of the fact] that’, ‘is *very* aware [of the fact] that’, ‘is *completely* aware [of the fact] that’, and ‘is *fully* aware [of the fact] that’ are semantically unproblematic and examples of them are easily found in English-language corpus searches. This difference in gradability has metaphysical bite. For often if we have two general conditions that we want to reduce, and one is gradable while the other not, *the direction of reduction involves taking the gradable condition as the more basic condition not the non-gradable condition*. This prompts the question: can at least some factive stative adjectives (e.g. ‘is aware that’) refer to states that are more basic than the states referred to by the non-gradable factive stative verbs (e.g. ‘knows that’)? And if so, can we also understand the remaining factive stative verbs (e.g. ‘understands that’) in terms of a factive stative adjective? Below I return to this issue of gradability *as well as the limited significance of English-language semantic intuitions and ways the intended external research group will help prevent this from being an anglo-centric project in epistemology*.

The final kind of argument against Knowledge-First Reductionism explored in ASK is an **explanatory argument**. Williamson (2000) provided some evidence for his knowledge-first epistemology by noting its ability to explain the long history of failed attempts to provide a reductive analysis of knowledge. ASK points to an alternative explanation: there is an equally long history of failures to distinguish knowledge of facts from the awareness of facts. Once this is done a new range of ways of analyzing knowledge and other epistemic notions in terms of factive adjectives become available—and have been shown to be promising (see ASK Ch.6-9, and Silva and Siscoe forth, *Noûs*). To whatever extent knowledge-first epistemology should be motivated and supported by Williamson’s observation about the history of failures to reductively analyze knowledge, awareness-first epistemology of the sort promoted by ASK also enjoys the same motivation and support. More than that, it is far from clear that knowledge-first epistemology has come through on its explanatory ambitions. For example, the strong majority, possibly all, knowledge-first epistemologists think that justified belief is to be somehow understood in terms of knowledge. But in the twenty three years since the appearance of *Knowledge and Its Limits* there has been at least 14 different knowledge-first theories of justification that have been pursued. All of which have been argued to be deeply problematic.¹⁹

3. Research objectives

So there are a range of arguments against Knowledge-First Reductionism and they converge on the idea that the factive adjectives track a mind-world relation that is distinct from the mind-world relation implicated by knowledge. So while the knowledge-first approach to questions in epistemology remains important, interesting, and well-worth continued study, the time is also ripe to begin exploring novel alternative approaches to foundational questions in epistemology that are neither traditionalist nor knowledge-first. This is a project that aims to achieve this. In particular,

¹⁸ Exploring Hetherington’s non-semantic arguments for the gradability of knowledge and how they relate to the gradability of awareness will be part of this project.

¹⁹ See Silva (2020) for a survey of the literature on knowledge-first theories of justification.



this project aims to engage in significant, cutting edge research by exploring a network of interrelated questions from new angles. There will be an emphasis on:

Subproject (1): exploring how factive adjectives and verbs – and the mental states they implicate – relate to each other,

Subproject (2): exploring how a better understanding of the factive adjectives can shed new light on debates about the nature of know-how,

Subproject (3): exploring how the mental states that the factive adjectives refer to might be normative in ways that knowledge is not, and

Subproject (4): exploring how a better understanding of the topics (1)-(3) can impact our understanding of internet epistemology, AI epistemology, collective epistemology, political epistemology, and the epistemology of encroachment.

The Work Programme. The work programme over the course of this project will consist of four subprojects. Each subproject will systemically contribute to an increasing understanding of the nature and place of the factives in contemporary epistemology and applications of this improved understanding to specific issues of importance within the field.

SUBPROJECT 1. FACTIVE ADJECTIVES: GRADABILITY AND RELATIONAL STRUCTURE. As noted above, there are a range of factive stative adjectives: ‘is aware that’, ‘is conscious that’, ‘is obvious to one that’, and ‘is clear to one that’. Further, all of these expressions are syntactically gradable and thus purport to refer to some kind of property (or properties) that *come in degrees*. We refer to degrees of awareness (/consciousness, /ignorance, /unawareness) of facts with expressions such as: ‘is *somewhat* aware (/conscious, /ignorant, /unaware) of the fact that’, ‘is *vaguely* aware (/conscious, /ignorant, /unaware) of the fact that’, ‘is *very* aware (/conscious, /ignorant, /unaware) of the fact that’, ‘is *completely* aware (/conscious, /ignorant, /unaware) of the fact that’, and ‘is *fully* (/conscious, /ignorant, /unaware) aware of the fact that’. Similarly, we refer to degrees of obviousness and clarity with expressions such as: ‘p is *more obvious* [/clearer] to me than q’ or ‘p is *more obvious* [/clearer] to me than it is to her’. These are semantically unproblematic expressions and examples of them are easily found in corpus searches.²⁰

Let’s start with some questions about the gradability of awareness. Gradable claims about awareness sometimes seem to track one’s degree of attention to a fact, but they can also track one’s broadly evidential position in relation to a fact. I am, for example, *completely aware* of the fact that I exist. But I can be completely aware of this fact even if I’m not directing any attention to it or to any other fact that presupposes my existence. For another example, suppose that you were *informed of the fact* that I gave a lecture last week from a reliable person. But my students were physically present for the entire lecture, both *hearing and seeing me lecture* for over an hour. The testimony you received made you aware of the fact that I gave a lecture last week, but my students who

²⁰ Here are some examples involving factual awareness gathered from English Corpora.org (<https://www.english-corpora.org>): “We wake up every day, only *slightly aware* that it’s a different day from the day before it”; “Most people are *vaguely aware* of the fact that cannabis has existed in a weird limbo in California for the past 21 years – not quite legal and yet not quite illegal”; “Newer solvers probably don’t even notice the black square pattern in the grid, but I would bet that many of them are at least *somewhat aware* that the pattern is usually symmetrical”; “By and large households are *very aware of the fact* that there’s a near-term surge in inflation, so that they’re upgrading their one-year-ahead beliefs”; “There is a new group of consumers who are *completely aware* that climate change is the problem of our generation, and yet they’re running around in plastic leggings with plastic shoes”; “His actions suggest he is *fully aware* that the dominance of a resentment narrative serves the GOP’s interest.”



personally witnessed my lecture for more than an hour are *even more* aware of the fact that I gave a lecture. So while you are *aware* of the fact that I gave a lecture, my students are *even more aware* of that fact.

If these claims are taken at face value, it seems plausible that to explain such comparative claims we need not refer to degrees of attention, but to refer to differences in the evidential or, broadly epistemic, positions that agents can occupy in relation to a fact. Such *provisional* ideas raise various questions. Is awareness gradable along an evidential dimension as the previous example suggests? If so, how can we best characterize the evidential dimension along which awareness is gradable? How widely shared are these evidential gradability judgements? Are there both subjective and objective dimensions along which awareness is gradable? That is, does it make a difference to gradability judgments that, in addition to having great evidence, one has some kind of perspective on the quality of one's evidence? Can negative higher-order judgments to the effect that one's evidence is less reliable than one thought decrease one's degree of awareness? If so, does it depend on whether or not that evidence is misleading? Some gradable adjectives are *multidimensional*, i.e. an adjective whose application conditions depend on *more than one* degreed threshold being met. Is 'aware of' a multidimensional gradable adjective? If so, what are those dimensions beyond evidential dimensions? Why is it that degrees of attention can sometimes impact gradability judgments involving awareness when attention is not required for awareness?²¹ Could it be because attention might, in part, be doubly representational—involving a representation of a representation? If so, should degrees of attention really be rejected as a further dimension of gradability and treated rather as a second, higher-order state of awareness? And how does all of this relate to the idea of *being conscious of the fact that p* and the degreed character of such consciousness?

We will also investigate the best way of understanding the gradability of other factive stative adjectives. Comparative judgments involving the expressions 'is obvious that' and 'is clear that' are semantically felicitous. We can take two facts that are obvious, and rank one as *more obvious* [*clearer*] *than* the other. The fact that this research proposal is written in English is both clear and obvious, and it is *clearer and more obvious than* the fact that this research proposal was composed by a human being rather than ChatGPT or a Google AI. The gradability of these factive adjectives raises its own questions. Could a fact be clear or obvious even if it is neither clear nor obvious to anyone? If so, then should we resist the idea that 'it is clear that p' and 'it is obvious that p' is somehow closely related to factual awareness of p, which requires *a subject* who is aware of the fact that p? Or, alternatively, can we understand being clear or obvious in terms of awareness of facts? If so, should we think of the being clear and being obvious in connection with *the ease* with which *people* could become aware of them? But what is 'ease' and which 'people' matter? Should we cash out 'ease' in terms of nearby worlds (i.e. in counterfactual terms)? If so, this would come close to how many people understand 'being in a position to know', i.e. being in a position where one *would* come to have knowledge were they to respond to their current situation in suitable ways. This in turn raises the question: what is the relation between a fact being obvious or clear and one being in a position to know it? Alternatively, instead of understanding clarity and obviousness in

²¹ Factual awareness does not require attention. Since you memorized your multiplication table you are aware of the fact that the product of 5 and 5 is 25. But you were also aware of that fact even before I drew your attention to it just a moment ago; it is a fact that you were aware of but not one that you were directing any attention towards. For further discussion see Chalmers (1996: 221) and Silva (2023: Chapter 3).



terms of counterfactuals, should we seek to understand it in terms of *agential dispositions*?²² What advantages might such a view have?

Further, we tend to assume that facts that are clear or obvious are of special normative significance: they are the kind of facts that one is responsible for being sensitive to in action, assertion, and belief. But what explains that? Will our best explanation of the properties that explain the gradability of ‘is clear that’ and ‘is obvious that’ support the idea that, sometimes, degrees of responsibility can be impacted by degrees of obviousness or clarity?

I have been focusing on factive *stative* adjectives, i.e. factive adjectives which indicate that one stands in a factive representational state. But there are also factive *non-stative* adjectives: ‘is unaware of the fact that’, ‘is ignorant of the fact that’, and ‘is oblivious to the fact that’. ‘Unaware’ and ‘ignorant’ also come in degrees: one can be *somewhat* unaware (/ignorant) of the fact that p, *very* unaware (/ignorant) of the fact that p, and *completely* unaware (/ignorant) of the fact that p. Corpus searches quickly reveal a rich range of examples of such expressions. Questions to investigate here concern the relations the factive stative adjectives bear to these factive non-stative adjectives. Silva and Siscoe (2023, *Noûs*) have argued that we must understand ignorance in terms of the absence of awareness, rather than the absence of knowledge²³ or the absence of true belief.²⁴ But this investigation left various questions unanswered. Some recent views of ignorance hold that ignorance has some kind of normative component.²⁵ For example, to be ignorant of the fact that p it must be the case that one ought to have satisfied some relevant normative standard. One might think that to be ignorant of the fact that p one *should have known* that p. Strictly speaking, it is possible to endorse the idea that ignorance is a lack of awareness together with the idea that one should have known. But there are other ways of specifying the relevant normative dimension. This issue will be explored. Second, as noted above, ignorance comes in degrees and, semantically, ‘being oblivious to the fact that p’ seems to express a stronger idea than ‘being merely somewhat ignorant of the fact that p’ and arguably it expresses the same idea as ‘being completely ignorant of the fact that p’. But if there’s a difference between being *merely somewhat* ignorant of a fact and being *completely* ignorant of it, what grounds that difference? And does it have any normative significance? For example, might degrees of accountability be partially connected to degrees of ignorance?

SUBPROJECT 2. (ANTI-)INTELLECTUALISM & AWARENESS-HOW. Awareness (or at least awareness-talk) takes different objects. You can be aware of *particulars* (aware of the cat), you can be aware of *properties of particulars* (aware of the stillness of the cat), you can be aware of *facts* (aware of the fact that a cat is nearby), you can be aware of *qualia* (aware of what it’s like to see a cat), and you can be aware of *propositions* (aware of the meaning of the sentence ‘Some cats bite’).²⁶ But you can also be aware of *skills* (aware of how to scare cats). This last kind of awareness, **awareness-how**,

²² For well-known reasons it is widely thought that dispositions cannot be identified with counterfactuals despite being a modal property. See Wasserman (2007, 2008, 2011) and Vetter (2015).

²³ Blome-Tillman (2016: 96), Fields (1994: 403), Fine (2018: 4032), Haack (2001: 25), Le Morvan (2011a), (2011b; 2012; 2013; 2019), K. Lynch (2016: 509), Unger (1975: 93), Zimmerman (1988: 75; 2008: ix).

²⁴ Goldman (1986: 26) and (1999: 5), Goldman and Olsson (2009: 19–21), Guerrero (2007: 62–63), Peels (2010; 2011; 2012; 2014), and Van Woudenberg (2009: 375).

²⁵ Meylan (2022) and Pritchard (2021).

²⁶ It is possible that the expression ‘is aware of’ refers to no neatly unified relation across all of its uses involving different objects. In this subproject I explore the idea that ‘is aware of how’ and ‘is aware of the fact that’ are unified parallel to the way intellectualists seek to unify know-how and know-that. But this leaves open the idea that awareness of particulars and properties cannot be unified with either awareness-how or factual awareness. This will be investigated further in the research proposal.



is a matter of being *aware of how to do something*. Corpora searches provide a rich range of examples where the expression ‘is aware of how to Φ ’ is used.²⁷ So just as there is knowledge-that and knowledge-how, there is awareness-that (factual awareness) and awareness-how. This raises a range of as-yet unexplored questions.

Before turning to those questions it’s worth highlighting the theoretical importance of knowledge-how. First, know-how is explanatorily significant for our understanding of sophisticated biological organisms (dogs, cats, humans, etc.). We cannot explain the distinctive survival-related successes of such biological organisms without reference to the fact that *they know how to do things*: they know how to find food, they know how to avoid predation, and so forth. Further, as Carlotta Pavese (2021) has pointed out, psychologists and neuroscientists have explicitly appealed to Ryle’s distinction between know-how and know-that when interpreting their empirical findings (Cohen & Squire 1980; Anderson 1983). An additional reason know-how is important concerns our understanding of the unity of pedagogical norms. Turri and Buckwalter (2014) have argued that know-how and know-that are both normatively significant for instruction:

Knowledge Norm for Pedagogy. (Asserting) If you do not *know that* p , then you should not assert that p . (Showing) If you do not *know how* to Φ , then you should not provide instruction on how to Φ .

Lastly, Lord (2018: 123-24) has argued that know-how is central to understanding how people come to possess reasons for responses:

Knowledge Norm for Possession. Necessarily, one possesses the fact that p as a reason for Φ -ing if and only if (i) one is in a position to *know that* p , (ii) one *knows how to use* p to Φ , and (iii) one is in a position *to manifest that know-how*.²⁸

Now, when it comes to understanding the nature of know-how and its relation to knowing-that, there are two primary views that emerge in response to the following question:

Can we reduce knowing-how to a kind of knowing-that?

Intellectualists argue that know-how just is a kind of knowledge-that (Stanley & Williamson 2001; Stanley 2011; Pavese 2015, 2017). According to one expression of intellectualism: what it is for a subject to know how to Φ is for that subject to have propositional knowledge of a true answer to the question ‘How could she herself Φ ?’, where an answer to this practical question has the form ‘ w is a way she herself could Φ ,’ for some way w . *Anti-intellectualists* argue this is false: know-how cannot be reduced to propositional knowledge of ways of Φ -ing.

²⁷ English Corpora.org (<https://www.english-corpora.org>) provides the following examples: “The Irish Cancer Society (ICS) is encouraging women to be *aware of how to* reduce their risk of cervical cancer”; “But it seems that not many of them are *aware of how to* protect themselves and their partners against STDs”; “...prisons are *very aware of how to* exploit Facebook’s Terms of Service...”; “When people aren’t *aware of how to* balance their vegan diet properly, they often forgo adequate plant sources...”; “... tinnitus associated with noise exposure could be prevented if people were more informed and *aware of how to* look after their hearing.”

²⁸ Condition (i) of this principle was opposed ASK, Chapter 6. There I explained the benefits of a view on which *factual awareness*, rather than propositional knowledge, is what is needed for condition (i), i.e. the *epistemic requirement* for the possession of reasons. But ASK was neutral on, what Lord calls, *the practical requirement* for the possession of reasons, i.e. conditions (ii) and (iii). Lord argues that it is *know-how*. This subproject will, among other things, explore the question of whether it is awareness-how, rather than know-how, that gives us the correct practical requirement.



Anti-intellectualists have sought to justify their view by pointing out that propositional knowledge is, in various ways, distinct and more demanding than knowledge-how. Specifically:

Non-Overlapping Properties (NOP). Propositional knowledge is not gradable,²⁹ it requires belief, it requires undefeated justification, and it requires the absence of environmental luck.³⁰ In contrast, knowledge-how is gradable, and it does not require belief,³¹ or undefeated justification,³² or the absence of environmental luck.³³

NOP has been a highly influential consideration in favor of rejecting intellectualism and endorsing some kind of anti-intellectualism.

It is at this point that the investigation of awareness-how and factual awareness opens up new perspectives on the debate about know-how. As noted above, there is very good reason to think that factual awareness (=being aware of facts) is gradable and there is also very good reason to think that factual awareness does not require belief, undefeated justification, or the absence of environmental luck (ASK, Chapter 4). Accordingly, there is nothing like NOP to justify the rejection of the following thesis:

A-Intellectualism. Knowledge-how just is a kind of factual awareness.

A very rough and unrefined version of A-Intellectualism might be as follows: what it is for a subject to know how to Φ is for that subject to be *aware of the fact that w is a way she herself could Φ* , for some way *w*. This claim about know-how is, on the surface, not at all jeopardized by the truth of NOP. And if A-Intellectualism is true, then presumably both know-how and awareness-how are a kind of factual awareness.

This raises a complex network of questions. First, might it be that NOP is false because it is *awareness-how and not know-how* that is gradable and fails to require belief, undefeated justification, and the absence of environmental luck? Put differently, have theorists been mistaking properties of know-how with properties of awareness-how? If so, do *all* the proposed counterexamples to intellectualism about know-how simply mistake know-how for awareness-how? Or are there some counterexamples to intellectualism about know-how that survive even if we've carefully distinguished know-how from awareness-how. Second, if know-how is gradable can its gradability be explained by the gradability of factual awareness? Or is factual awareness gradable in ways that do not explain the gradability of knowledge-how? Third, if we can reduce both awareness-how and know-how to factual awareness, what form does that reduction take? And is awareness-how identical to know-how or is know-how in some way more demanding than awareness-how? Fourth, if awareness-how is in some ways distinct from know-how, which of these two 'how' relations is central to norms for pedagogy (cf. Buckwalter and Turri 2014) and the possession of reasons (cf. Lord 2018)? Should we, in other words, reject the above knowledge norms and instead prefer some kind of awareness norm in their place? This last question brings us to the next subproject.

²⁹ Ryle (1949), Dretske (1981: Ch. 5), Stanley (2005: Ch2), and Brogaard (2016: 57).

³⁰ For discussion and defense of these latter claims see ASK Chapter 2-4 and 8-9.

³¹ Dreyfus (1991, 2005), Wallis (2008), Brownstein & Michaelson (2016).

³² Carter & Navarro (2017).

³³ Cath (2011), Poston (2009), and Carter & Pritchard (2015).



SUBPROJECT 3. THE NORMATIVITY OF FACTUAL AWARENESS VS KNOWLEDGE. There are epistemic norms for assertion, action, and deliberation, i.e. there are epistemic conditions that must be satisfied if one is to acceptably assert that *p*, act on *p*, or treat *p* as a premise in deliberative reasoning. The idea that knowledge plays an essential role in these norms is supported by our social practices of epistemic accountability, advice giving, as well as broadly Moorean phenomena. If someone acts on *p*, asserts that *p*, or deliberates on the basis of *p*, it's appropriate to ask or challenge them with a knowledge-question: do you know that *p*? We take it to be correct to advise people to not act on *p* if they don't know that *p*.³⁴ And assertions of the form 'p, but I don't know it' are Moore-paradoxical in a way that calls out for explanation. As advocates of knowledge norms have observed, knowledge norms can provide an explanation of such phenomena.

These, together with many further considerations, have been argued to support the following norms:

(K-norm) One can assert *p*, act on *p*, or treat *p* as a premise in deliberative reasoning only if *one knows p (or is at least in a position to know p)*.^{35 36}

(JK-norm) One can assert *p*, act on *p*, or treat *p* as a premise in deliberative reasoning only if one has *undefeated justification to believe that one knows p (or is at least in a position to know p)*.³⁷

But if knowledge is distinct from factual awareness we must assess how awareness-centric versions of these norms would fare:

(A-norm) One can assert *p*, act on *p*, or treat *p* as a premise in deliberative reasoning only if one is *aware of the fact that p* (or else is *in a position to be aware of the fact that p*).

(JA-norm) One can assert *p*, act on *p*, or treat *p* as a premise in deliberative reasoning only if one has *undefeated justification to believe that one is aware of the fact that p (or is at least in a position to be aware of the fact that p)*.

Because knowledge that *p* entails awareness of the fact that *p*, the (K-norm) entails the (A-norm) and the (JK-norm) entails the (JA-norm). But the awareness norms differ from their knowledge-centric counterparts in virtue of awareness being less demanding than knowledge. This might give awareness-centric norms an edge in certain regards.³⁸

³⁴ Though one may act on their knowledge of *the probability of p*, or act on their knowledge *that acting on p is a dominant course of action*, or act on their knowledge *that p is close to true (i.e. if false, not false in a way that makes a practical difference)*, or act on their knowledge of other distinct *p*-related propositions.

³⁵ DeRose (2002), Reynolds (2002), Adler (2002: 275), Hawthorne (2004), Stanley (2005), Stanley & Hawthorne (2008), Engel (2008), Schaffer (2008), and Turri (2011).

³⁶ There is also an alternative justification norm: one can assert *p*, act on *p*, or treat *p* as a premise in deliberative reasoning only if one has undefeated justification to believe *p*. For defense of norms of this sort see Douven (2006, 2009), Lackey (2007, 2008), Gerken (2012, 2014, 2017), McKinnon (2013, 2015), and Reuter & Brössel (2018).

³⁷ Smithies (2012) and Reynolds (2013).

³⁸ For example, the (A-norm) seems better suited to handling anti-luck problems that the (K-norm) faces. Smithies argues that the (K-norm) is problematic as it forbids assertion, action, and deliberation on *p* in cases of environmental luck. But there are reasons for thinking awareness is compatible with environmental luck. See ASK Chapter 4 for discussion.



There are many issues to consider in regard to the debate between awareness-centric and knowledge-centric norms. These norms assert a single unified norm for assertion, *and* action, *and* deliberation. But could there be different norms for these kinds of activities? If so, would they all be knowledge-centric or could some be awareness-centric? Given that awareness comes in degrees, what role might degrees of awareness play in understanding the normativity of awareness? Further, what relation do these norms have to norms for our class of doxastic states? The set of distinct doxastic attitudes involves outright belief, credence, faith, expectation, presupposition, and suspension. Should norms for ALL of these states be knowledge-centric, or should some or all of them be awareness-centric? For example, take belief. Some have argued that the belief is weak (compatible with lacking a very high credence) and thus is normatively less demanding than assertion (Hawthorne, et al. 2016). If true, should we prefer an awareness-centric norm for belief, while preferring a knowledge-centric norm for assertion? Or should both get an awareness norm, where the degree of awareness required for appropriate assertion is less than the degree of awareness required for appropriate belief? Further, could the norms for action, assertion, and deliberation as well as for the range of our doxastic states be in any way context sensitive or sensitive to practical or moral stakes? Further, what role, if any, should genealogical considerations play in determining the relevant norms? For example, suppose our concept of knowledge is connected with our need to trust other people as informants.³⁹ After all, we cannot be everywhere at once and we need to rely on other people for information about the world that we cannot easily acquire for ourselves. Suppose, then, that the expression ‘S knows that p’ came about to serve the function of identifying S as trustworthy when it comes to inquiring of S whether or not p is true.⁴⁰ If so, might that provide a significant reason to endorse a knowledge-norm of assertion? Furthermore, could genealogical considerations also be leveraged to support other knowledge-centric norms for action, deliberation, and our other doxastic states—thereby limiting the normative role of states of awareness that fail to amount to knowledge?

SUBPROJECT 4. NEW DIRECTIONS IN EPISTEMOLOGY. Contemporary epistemology has turned to addressing a range of issues that were either absent or much less prominent in pre-2000 epistemology. This final subproject will explore how the results of subprojects 1-4 impact these new directions in epistemology.

Internet Epistemology. The epistemology of learning has direct consequences for how we think about the epistemology of the internet and our on-line epistemic practices. Consider, for example, that we *learn* a lot on-line by googling and *responsibly* searching through the results. But even when we find a reliable source of information *that gives us the truth* about whether p, we often are in a position where *if we had done just a little bit more searching* (say scrolling through one or two more pages of search hits) we could easily have found additional and sufficiently credible defeating information which would have (rationally) led us not to believe p. Certain anti-luck epistemologies and certain theories about the defeating power of higher-order evidence imply that the ease of access to defeaters obstructs knowledge-acquisition (cf. Ballantyne 2015, 2019), and therefore obstructs learning. This gives rise to a dilemma: either epistemologies that have this implication are false or we don’t *learn* very much on-line. But if *learning* is more closely associated with factual awareness, and if factual awareness does not involve as robust an anti-luck element as

³⁹ Craig (1990), Fricker (2008), Pritchard (2012a), Queloz (2021).

⁴⁰ For such considerations see Craig (1990), Fricker (2007), Pritchard (2012a), and Hannon (2019).



knowledge and/or is consistent with negative higher-order information, then new responses to this dilemma begin to emerge.

There are likewise questions about the extent to which gathering information in echo chambers and information bubbles impacts the possibility of knowledge-acquisition *even when one is ignorant that one is evidence-gathering in such an on-line environment*.⁴¹ For example, suppose that one gathers information in an environment that is disposed to resist climate science, and produces reasons against trusting sources/arguments that indicate, for example, that disastrous climate change is anthropogenic. Depending on various further factors, this could ensure a certain degree of *unreliability in inquiry*. What impact might this have on our ability to learn and know in such environments even when the on-line environment is producing good reasons for *other* climate-related conclusions? Could we become aware of facts without knowing them in such on-line spaces due to environmental biases present? What conditions have to obtain? Could the epistemic impact be one of degree, and thus the kind of impact only appreciable with a proper understanding of the factive adjectives, which allow us to assess one as having *partial* awareness or *partial* ignorance? All of this raises many previously unexplored questions about the nature of learning and the epistemological value of our on-line epistemic practices. For obvious reasons, having a well-worked-out epistemology of our on-line epistemic practices has become a new and quickly growing area of research for philosophers,⁴² and these issues will be engaged.

AI Epistemology. The creation of Large Language Models (LLMs) are an unprecedented advance in recent technology. OpenAI, claims to have

trained a [LLM] called ChatGPT which interacts in a conversational way. The dialogue format makes it possible for ChatGPT to answer followup questions, admit its mistakes, challenge incorrect premises, and reject inappropriate requests.

This description of ChatGPT is, like other descriptions of this LLM from the technology sector, riddled with robust cognitive-epistemic language. ChatGPT is said to have been ‘trained’, which implies *learning*, and therefore, a transition from *ignorance* to *knowledge*. It is said to be able to ‘answer follow up questions’, which implies *knowledge of answers* or at least *knowledge of possible answers*. It is said to ‘admit mistakes’, which implies *knowledge of past errors*. It is said to ‘reject inappropriate requests’, which implies the ability to *recognize deviations from a standard*. Elsewhere, ChatGPT is said to ‘hallucinate’ when it asserts falsehoods that are easily recognized as such by human agents, this suggests the ability to *mistakenly treat itself as knowing*.

When it comes to the cognitive-epistemic language applied to LLMs, there appear to be two interpretations. The preferred interpretation in the technology sector and elsewhere is *anthropomorphic*: cognitive-epistemic attributions to LLMs do not strictly or literally apply to any LLM. LLMs do not really know (learn, remember, etc.) things because *they are not agents with genuine mental states*; rather, such language is, at most, aptly used only as metaphor. An alternative interpretation is *prospectively literal*: while current LLMs may fail to literally know (learn, remember, etc.), *sufficiently matured* LLMs will be agents with the capacity to literally know (learn, remember, etc.) things in the very same sense that human beings do. As already indicated, the prospectively literal interpretation is beyond the credulity of many and the underlying rationale for

⁴¹ Elzinga (2020), Nguyen (2020a, 2020b), and Lackey (2021).

⁴² Lynch (2016), Gunn and Lynch (2019), Coady and Chase (2018), and Fricker, et al. (2019).



this is connected to the fact that it seems unlikely that LLMs are *conscious agents*, and epistemic language seems to be literally applicable only to *conscious agents*.

But there is a third view, one that paves a middle path. It involves the idea that there is *some* cognitive-epistemic language that can literally apply to sufficiently matured LLMs because such LLMs can instantiate *some* cognitive-epistemic properties, and they can instantiate such properties *whether or not they are conscious* (cf. Chalmers 1996). Call this the *emergent* interpretation of LLMs. The foundation of the emergent interpretation rests on the idea that consciousness is not necessary for the instantiation of at least some robust cognitive-epistemic properties. In particular, recent work on the metaphysics of awareness suggests that awareness is sufficiently distinct from consciousness that it allows us to begin to understand how sufficiently matured LLMs can have the capacity to be fully aware of things in the very same sense that humans are able to be fully aware of things. Against this, however, some argue that states as robust as propositional knowledge cannot be instantiated by non-conscious systems (Smithies 2019).

This sub-project will be guided by the following general questions:

1. How is *knowledge* related to both awareness-of and consciousness-of?
2. Genuine agency requires the possession of and a responsiveness to reasons. If non-conscious beings can be in states of awareness-of, can they also possess and be responsive to reasons? What range of responses are available for such non-conscious beings?
3. How might we understand the intellectual virtues in relation to agents which have the capacity to be aware, but fail to have the capacity to be conscious?
4. As people begin to rely on LLMs for new information (knowledge-acquisition), how do answers to the previous questions impact our understanding of knowledge-acquisition processes that involve (apparent) testimony from LLMs?

Political Epistemology. Political agents (chancellors, presidents, senators, mayors, etc.) have to manage a wide range of affairs on issues where the best course of action depends on facts that are beyond their expertise and understanding. Further, political agents have to make decisions that impact their constituents, and they must justify their decisions to constituents who often have even less understanding of the relevant affairs than they do. Finally, political agents have to make their decisions in contexts of significant (apparent) peer disagreement, which would seem to non-trivially impact the degree of justification they have for their assertions and their actions. As noted above, many epistemologists believe that there is a growing body of evidence that our justificatory practices for action and assertion are governed by knowledge norms. But if we assess political agents relative to knowledge norms, they will typically fall short of any positive assessment by these knowledge norms. Yet this should seem strange, for we often witness imperfect political agents doing well, and succeeding by some normatively significant standard to make justifiable decisions. But it would be strange to assess politicians as so often doing well when they regularly fail to meet the knowledge-centric criteria for making justified and justifiable decisions. So there is a dilemma between the frequent inability of political agents to satisfy knowledge norms and the idea that political agents do not infrequently do well in terms of making and justifying their decisions. We will explore the question of whether or not we should think of “good” or “virtuous” political agency in terms of awareness or knowledge. For the less demanding character of factual awareness allows us to assess political agents against a less demanding standard. Some relevant works



associated with this topic include Baghranian and Croce (2021), Boulton (2021), and Hannon (2022).

Collective Epistemology. Collective epistemology is concerned with questions about the knowledge and rationality of *groups of agents* as well as *group agents* (=collective agents). Collective epistemology has been a somewhat less prominent area of epistemology, but it's been on the rise in the last 20 years and has quickly become a well-defined, distinctive area of epistemological inquiry (Haddock, Millar, Pritchard 2010; Goldman and Whitcomb 2011; Fricker, et al. 2019; Lackey 2020). A significant aspect of collective epistemology concerns the rationality and knowledgeability of *group attitudes*: group belief, group rationality, group knowledge. At present, there is no systematic discussion about the nature of awareness for groups—at least there is no discussion of this that doesn't tend to assume that group awareness just is group knowledge. This project will investigate general principles for group knowledge and group rationality, and how they relate to group awareness. There is good reason to expect unique challenges here. For principles that explain the knowledge and rationality of individuals have not easily generalized to principles that explain the knowledge and rationality of groups (e.g. Lackey 2016; Silva 2018a). So it's reasonable to expect that similar complexities will arise in the case of group awareness. The project will seek to navigate these complexities in order to arrive at a theoretically viable view about the extent to which the factive states referred to with the factive adjectives can apply to groups.

Encroachment (Moral and Pragmatic). Many have defended the idea that knowledge is subject to pragmatic encroachment. That is, knowing *p* can depend on practical factors, e.g. suppose one would die or lose their home were one to act on *p* when *p* is false (Stanley 2005; Fantl and McGrath 2009; DeRose 2009). Moreover, some have also argued that knowledge is subject to moral encroachment, e.g. that one cannot know someone is guilty on the basis of their race or gender even if their race or gender is statistically highly predictive of guilt and one is in fact guilty. Why? Because forming beliefs in these ways—even if reliably true in certain contexts—is nevertheless morally problematic (Schroeder and Basu 2019; Basu 2019; Kim and McGrath 2019). Contemporary epistemologists are divided on pragmatic encroachment and moral encroachment. But once knowledge is separated from awareness new theoretical options arise. For example, it *may* be that knowing really is determined in part by practical or even moral factors, while the awareness of facts is not. In which case, it would be possible to adjudicate disputes about pragmatic and moral encroachment in an entirely new way. As of yet there is no exploration of this approach to debates about encroachment in the growing literature on moral and pragmatic encroachment.

4. Methodology

Research Methods. The method of conceptual analysis, the method of conceptual engineering, the method of genealogical inquiry, and the method of phenomenal generalization, will be employed alongside the method of seeking reflective equilibrium. While conceptual analysis seeks to analyze concepts by considering our intuitions about the applicability of terms/concepts in concrete cases, genealogical inquiry asks about what practical pressures might have led us to form and rely on certain concepts. The underlying idea is that we can better understand our concepts by better understanding the pressures that led us to develop them (cf. Craig 1990; Hannon 2019; Queloz 2021). Conceptual engineering, in contrast, asks questions about whether our concepts (or words) are in need of revision, owing to a failure of our concepts (words) to adequately track real-world relations (cf. Cappelen 2018). As this project aims to understand the class of *mental*



states that factive statives are used to refer to, as is standard in philosophical investigation, we will also make use of our first-personal acquaintance with our own mental states and dispositions to make mental state attributions to help us better understand what kinds of states the factives refer to. For just as our *specific experiences* of pain and of perception helps us understand what kinds of states that the terms ‘pain’ and ‘perception’ refer to *in general*, our specific experiences of being in factive states or inclinations to make certain attributions can sometimes help us understand what kinds of states the factives refer to in general. We will also be studying the logical relations and interaction of possible general principles governing the factives and other epistemic relations and applying the method of reflective equilibrium to our assessment of potential normative principles involving factives.⁴³ These philosophical investigations will be informed by relevant literature from linguistics and psychology.

References

- Adler, J. E. (2002). *Belief's Own Ethics*. Cambridge: MIT Press.
- Anderson, John R., 1983, *The Architecture of Cognition*, Cambridge: Harvard University Press.
- Audi, R. (1998). *Epistemology: A contemporary introduction to the theory of knowledge*. Routledge.
- Ballantyne, N. (2015). The Significance of Unpossessed Evidence. *Philosophical Quarterly*, 65(26), 315-335.
- Ballantyne, N. (2019). *Knowing Our Limits*. Oxford: Oxford University Press.
- Basu, R. (2019). What We Epistemically Owe To Each Other. *Philosophical Studies*.
- Bengson, J. 2015. The Intellectual Given. *Mind*, 124(495), 707-760.
- Bernecker, S. (2009). *Memory: A Philosophical Study*. Oxford: Oxford University Press.
- Bird, A. (2007). Justified Judging. *Philosophy and Phenomenological Research*, 74(1), 81-110.
- Bonjour, L. (1985). *The Structure of Empirical Knowledge*. Cambridge: Harvard University Press.
- Baghrarian, Maria & Croce, Michel (2021). Experts, Public Policy and the Question of Trust. In Michael Hannon & Jeroen De Ridder (eds.), *Routledge Handbook of Political Epistemology*. London: Routledge.
- Blome-Tillmann, Michael (2016). Ignorance and Epistemic Contextualism. In Rik Peels and Martijn Blaauw, eds., *The Epistemic Dimensions of Ignorance*. Cambridge: Cambridge University Press, 96-113.
- Brogaard, B. (2016). Ignorance and Incompetence: Linguistic Considerations. In *The Epistemic Dimensions of Ignorance*. Edited by Rik Peels and Martijn Blaauw. Cambridge University Press.
- Brownstein, M. & E. Michaelson. (2016). Doing without Believing: Intellectualism, Knowledge-How, and Belief-Attribution. *Synthese*, 193(9), 2815–2836.
- Boult, Cameron (2021). The (Virtue) Epistemology of Political Ignorance. *American Philosophical Quarterly*.
- Buckwalter, W. and J. Turri. (2014). Telling, showing and knowing: a unified theory of pedagogical norms. *Analysis*, 74(1)16-20.
- Carter, J. A. & J. Navarro (2017). The Defeasibility of Knowledge-How. *Philosophy and Phenomenological Research*, 95(3): 662–685.
- Carter, J. A. & Duncan Pritchard (2015). Knowledge-How and Epistemic Value. *Australasian Journal of Philosophy*, 93(4), 799–816.
- Carter, J. A., E. Gordon, & B. Jarvis, eds. (2017). *Knowledge First: Approaches in Epistemology and Mind*. Oxford: Oxford University Press.
- Cath, Y. (2011). Knowing How Without Knowing That. In Bengson and Moffett 2011: 113–135.
- Cappelen, H. (2018). *Fixing Language: An Essay on Conceptual Engineering*. Oxford: Oxford University Press.

⁴³ This method of reflective equilibrium is widely used in normative philosophy. It consists in working back and forth among our (i) considered judgments about particular scenarios, (ii) the implications of general rules to those same scenarios, and (iii) an assessment of the consistency and relation between the (i) and (ii). To the extent (i) and (ii) are incompatible or at odds, we use the value we associate with these judgments and principles to anchor either a revision of the general principles or a rejection of the earlier considered judgements. Revisions to general principles are subsequently tested in application to further concrete hypothetical cases, and depending on the acceptability of their implications, they are revised until we have reached an optimal coherence. For more on this see method see Pust (2021: Ch.1).



- Carruthers, P. & R. Gennaro (2020). Higher-Order Theories of Consciousness. The Stanford Encyclopedia of Philosophy (Fall 2020 Edition), Edward N. Zalta (ed.), URL = <<https://plato.stanford.edu/archives/fall2020/entries/consciousness-higher/>>.
- Chalmers, D. (1996). *The Conscious Mind*. Oxford: Oxford University Press.
- Coady, D. and J. Chase. (2018). *Routledge Handbook of Applied Epistemology*. New York: Routledge.
- Cohen, N. J. and L. R. Squire. (1980). Preserved Learning and Retention of Pattern-Analyzing Skill in Amnesia: Dissociation of Knowing How and Knowing That. *Science*, 210(4466): 207–210. doi:10.1126/science.7414331
- Craig, E. (1990). *Knowledge and the State of Nature: An Essay in Conceptual Synthesis*. Oxford: OUP.
- De Rose, K. (2002). Assertion, Knowledge, and Context. *The Philosophical Review*, 111(2), 167–203.
- De Rose, K. (2009). *The Case for Contextualism*. Oxford: Oxford University Press.
- Douven, I. (2006). Assertion, Knowledge, and Rational Credibility. *The Philosophical Review*, 115: 449–485.
- Douven, I. (2009). Assertion, Moore, and Bayes. *Philosophical Studies*, 144(3): 361–375.
- Dretske, F. (1981). *Knowledge and the Flow of Information*. Cambridge: MIT Press.
- Dretske, F. (1993). Conscious Experience. *Mind*, 102(406), 263–283.
- Dreyfus, H. L. (1991). *Being-in-the-World: A Commentary on Heidegger's Being and Time, Division I*, Cambridge: MIT.
- Dreyfus, Hubert L. (2005). Overcoming the Myth of the Mental: How Philosophers Can Profit from the Phenomenology of Everyday Expertise. *Proceedings and Addresses of the American Philosophical Association*, 79(2), 47–65.
- Edenberg, Elizabeth (2021). Political Disagreement: Epistemic or Civic Peers? In Michael Hannon & Jeroen De Ridder (eds.), *Routledge Handbook of Political Epistemology*.
- Elgin, C. (2006). From Knowledge to Understanding. In *Epistemology Futures*, ed. Stephen Hetherington Oxford: Oxford University Press, 199–215.
- Elzinga, B. (2020). Echo Chambers and Audio Signal Processing. *Episteme*: 1–21.
- Engel, P. (2008). In What Sense Is Knowledge the Norm of Assertion? *Grazer Philosophische Studien*, 77, 45–59.
- Fantl, J. and M. McGrath. (2009). *Knowledge in an Uncertain World*. Oxford: Oxford University Press.
- Fields, L. (1994). Moral beliefs and Blameworthiness. *Philosophy*, 69(270), 397–415.
- Fine, K. (2018). Ignorance of Ignorance. *Synthese*, 195(9), 4031–4045.
- Fricker, M. (2007). *Epistemic Injustice*. Oxford: Oxford University Press.
- Fricker, M. (2008). Scepticism and the Genealogy of Knowledge: Situating Epistemology in Time. *Philosophical Papers* 37 (1): 27–50.
- Fricker, M., P. Graham, D. Henderson, J. Wyatt, and N. Pedersen. (2019). *The Routledge Handbook of Social Epistemology*. New York: Routledge.
- Gerken, M. (2012). Discursive Justification and Skepticism. *Synthese*, 189(2), 373–394.
- Gerken, M. (2014). Same, Same but Different: The Epistemic Norms of Assertion, Action and Practical Reasoning. *Philosophical Studies*, 168(3), 725–744.
- Gerken, M. (2017). *On Folk Epistemology: How We Think and Talk about Knowledge*. Oxford University Press.
- Guerrero, A. (2007). Don't Know, Don't Kill: Moral Ignorance, Culpability, and Caution. *Philosophical Studies*, 136, 59–97.
- Goldman, A. (1986). *Epistemology and Cognition*. Cambridge: Harvard University Press.
- Goldman, A. (1999). *Knowledge in a Social World*. Oxford: Oxford University Press.
- Goldman, A. and E. Olsson (2009). Reliabilism and the Value of Knowledge. In Adrian Haddock, Alan Millar, and Duncan Pritchard, eds., *Epistemic Value*. Oxford: Oxford University Press.
- Goldman, A. and D. Whitcomb. (2011). *Social Epistemology: Essential Readings*. Oxford University Press.
- Greenough, P., & D. Pritchard, eds. (2009). *Williamson on Knowledge*. Oxford: Oxford University Press.
- Gunn, H. and M. Lynch. (2019). Googling. In D. Coady and J. Chase (Eds.), *Routledge Handbook of Applied Epistemology* (pp. 41–53). New York: Routledge.
- Haack, Susan (2001). 'The Ethics of Belief' Reconsidered. In Matthias Steup, ed., *Knowledge, Truth, and Duty: Essays on Epistemic Justification, Responsibility, and Virtue*. Oxford University Press, 21–33.
- Hannon, M. (2019). *What's the Point of Knowledge*. Oxford: Oxford University Press.
- Hannon, M. (2021). Recent Work in the Epistemology of Understanding. *American Philosophical Quarterly* 58(3): 269–290.



- Hawthorne, J. (2004). *Knowledge and Lotteries*. Oxford: Oxford University Press.
- Hawthorne, J. and J. Stanley. (2008). Knowledge and Action. *The Journal of Philosophy*, 105(10), 571-590.
- Hawthorne, John ; Rothschild, Daniel & Spectre, Levi (2016). Belief is weak. *Philosophical Studies* 173 (5):1393-1404.
- Haddock, A., A. Millar, and D. Pritchard. (2010). *Social Epistemology*. Oxford: Oxford University Press.
- Hannon, M. (2022). Are Knowledgeable Voters Better? *Politics, Philosophy, and Economics*, 21(1), 29-54.
- Hetherington, S. (2001). *Good Knowledge, Bad Knowledge*. Oxford: Oxford University Press.
- Huemer, M. (2001). *Skepticism and the Veil of Perception*. Lanham: Rowman & Littlefield.
- Hyman, J. (2017). II—Knowledge and Belief. *Aristotelian Society Supplementary Volume* 91 (1):267-288.
- Kelp, C. (2013). Knowledge: The Safe-Apt View. *Australasian Journal of Philosophy*, 91(2), 265-278.
- Kim, B. and M. McGrath. (2019). *Pragmatic Encroachment in Epistemology*. New York: Routledge.
- Kusch, M. 2013. 'Naturalized Epistemology and the Genealogy of Knowledge'. In *Contemporary Perspectives on Early Modern Philosophy: Nature and Norms in Thought*. Edited by Martin Lenz and Anik Waldow, 87–100. Springer.
- Lackey, J. (2007). Norms of Assertion. *Noûs*, 41(4), 594-626.
- Lackey, J. (2008). Norms of Assertion and Testimonial Knowledge. In her *Learning from Words: Testimony as a Source of Knowledge*, Oxford: Oxford University Press, pp. 103–141.
- Lackey, J. (2016). What is Justified Group Belief? *The Philosophical Review*, 125(3), 341-396.
- Lackey, J. (2020). *The Epistemology of Groups*. Oxford: Oxford University Press.
- Lackey, J. (2021). Echo Chambers, Fake News, and Social Epistemology. In Sven Bernecker, Amy K. Flowerree & Thomas Grundmann (eds.), *The Epistemology of Fake News*. Oxford University Press.
- Lehrer, K. (1974). *Knowledge*. Oxford: Oxford University Press.
- Le Morvan, P. (2011a). Knowledge, Ignorance, and True Belief. *Theoria* 77(1): 309-318.
- Le Morvan, P. (2011b). On Ignorance: A Reply to Peels. *Philosophia* 39(2): 335-344.
- Le Morvan, P. (2012). On Ignorance: A Vindication of the Standard View. *Philosophia* 40(2): 379-393.
- Le Morvan, P. (2013). Why the Standard View of Ignorance Prevails. *Philosophia* 41(1): 239-256.
- Le Morvan, Pierre and Rik Peels (2016). The Nature of Ignorance: Two views. In Rik Peels and Martijn Blaauw, eds., *The Epistemic Dimensions of Ignorance*. Cambridge: Cambridge University Press, 12-32.
- Littlejohn, C. (2015). Knowledge and Awareness. *Analysis* 75(4): 596-603.
- Littlejohn, C. (2022). *Knowledge-First Epistemology*. Oxford Bibliographies. DOI: 10.1093/OBO/9780195396577-0380
- Lord, E. (2018). *The Importance of Being Rational*. Oxford: Oxford University Press.
- Lynch, K. (2016). Willful Ignorance and Self-Deception. *Philosophical Studies*, 173(2), 505-523.
- Lynch, M. (2016). *The Internet of Us: Knowing More While Understanding Less*. New York: Norton & Co.
- Manley, David & Wasserman, Ryan (2007). A gradable approach to dispositions. *Philosophical Quarterly* 57 (226):68–75.
- Manley, David & Wasserman, Ryan (2008). On linking dispositions and conditionals. *Mind* 117 (465):59-84.
- Manley, D. & Wasserman, R. (2011). Dispositions, Conditionals, and Counterexamples. *Mind* 120 (480):1191-1227.
- McKinnon, R. (2013). The Supportive Reasons Norm of Assertion. *American Philosophical Quarterly*, 50(2): 121–135.
- McKinnon, R. (2015) *The Norms of Assertion: Truth, Lies, and Warrant*. London: Palgrave Macmillan UK.
- McDowell, J. (2002). Responses. In Nicholas H. Smith, ed., *Reading McDowell on Mind and World*. London: Routledge, 269–305.
- Meylan, A. (2022). In Defence of the Normative Account of Ignorance. *Erkenntnis*: 1-15.
- Mitova, V., ed. (2018). *The Factive Turn in Epistemology*. Cambridge: Cambridge University Press.
- Mulligan, K. and F. Correia. (2021). Facts. *The Stanford Encyclopedia of Philosophy* (Winter 2021 Edition), Edward N. Zalta (ed.), URL = <<https://plato.stanford.edu/archives/win2021/entries/facts/>>.
- Nagel, J. (2013). Knowledge as a mental state. *Oxford Studies in Epistemology* 4: 275–310.
- Nagel, J. (2017). Factive and nonfactive mental state attribution. *Mind and Language* 32(5): 525-44.
- Nguyen, C. Thi (2020a). Cognitive islands and runaway echo chambers: problems for epistemic dependence on experts. *Synthese* 197 (7):2803-2821.
- Nguyen, C. Thi (2020b). Echo chambers and epistemic bubbles. *Episteme* 17 (2):141-161.
- Paul, E.S. (2020). Cartesian Clarity. *Philosophers' Imprint*, 20(19), 1-28.
- Pavese, C. (2015). Practical Senses. *Philosophers' Imprint*, 15(29), 1-25.



- Pavese, C. (2017). Know-How and Gradability. *Philosophical Review*, 126(3), 345-383.
- Pavese, C. (2021). Pavese, Carlotta, "Knowledge How", *The Stanford Encyclopedia of Philosophy* (Fall 2022 Edition), Edward N. Zalta & Uri Nodelman (eds.), URL = <<https://plato.stanford.edu/archives/fall2022/entries/knowledge-how/>>.
- Peels, R. (2010). What is Ignorance? *Philosophia* 38(1): pp. 57-67.
- Peels, R. (2011). Ignorance is Lack of True Belief: A Rejoinder to Le Morvan. *Philosophia* 39(2): 344-355.
- Peels, R. (2012). The New View on Ignorance Undefeated. *Philosophia* 40(4): 741-750.
- Peels, R. (2014). What Kind of Ignorance Excuses? *The Philosophical Quarterly* 64(256): 478-496.
- Peels, R. (2020). Asserting Ignorance. In Sanford Goldberg, ed., *The Oxford Handbook of Assertion*. Oxford: Oxford University Press, pp. 605-624
- Phillips, Jonathan, Wesley Buckwalter, Fiery Cushman, Ori Friedman, Alia Martin, John Turri, Laurie Santos, and Joshua Knobe (2021). Knowledge before belief. *Behavioral and Brain Sciences* 44: 1-17.
- Poston, T.(2009).Know How to Be Gettiered? *Philosophy and Phenomenological Research*,79(3), 743-747.
- Pritchard, D. (2012a). Anti-Luck Virtue Epistemology. *The Journal of Philosophy* 109(3), 247-79.
- Pritchard, D. (2012b). *Epistemological Disjunctivism*. Oxford: Oxford University Press.
- Pritchard, D. (2021). Ignorance and Normativity. *Philosophical Topics*, 49(2), 225-243.
- Pust, J. (2021). *Intuitions as Evidence*. New York: Routledge.
- Queloz, M. (2021). *The Practical Origin of Ideas*. Oxford University Press.
- Reuter, K. & Brössel, P. (2018). No knowledge required. *Episteme*, 16(3), 303-321.
- Rotstein, Carmen & Winter, Yoav (2004). Total Adjectives vs. Partial Adjectives: Scale Structure and Higher-Order Modifiers. *Natural Language Semantics* 12 (3):259-288.
- Reynolds, S. (2002). Testimony, Knowledge, and Epistemic Goals. *Philosophical Studies*, 110(2), 139-161.
- Reynolds, S. (2013). Justification as the Appearance of Knowledge. *Philosophical Studies*, 163, 367-383.
- Ryle, G. (1949). *The Concept of Mind*. Chicago: University of Chicago Press.
- Schaffer, J. (2008). Knowledge in the Image of Assertion. *Philosophical Issues*, 18, 1-19.
- Schellenberg, S. (2018). *The Unity of Perception*. Oxford: Oxford University Press.
- Schroeder, M. (2015). Knowledge is Not the Most General Factive Stative Attitude. University of Southern California. <http://static1.squarespace.com/static/55505fc8e4b032b4451e4a90/t/55ca36f6e4b0423ff5ac5e96/1439315702607/Schroeder_Factive_Stative_Attitudes.pdf>.
- Schroeder, M. (2021). *Reasons First*. Oxford: Oxford University Press.
- Schroeder, M. and R. Basu. (2019). Doxastic Wrongoing. In B. Kim and M. McGrath (Eds.), *Pragmatic Encroachment in Epistemology* (pp. 181-205). New York: Routledge.
- Silva, P. and W. Siscoe. (forth). Ignorance and Awareness. *Noûs*.
- Silva, P. (forth). Self-Fulfilling Beliefs: A Defense. *Australasian Journal of Philosophy*.
- Silva, P. (2023). *Awareness and the Substructure of Knowledge*. Oxford: Oxford University Press.
- Silva, P. (2022). Basic knowledge and the normativity of knowledge: The awareness-first solution. *Philosophy and Phenomenological Research*. <https://doi.org/10.1111/phpr.12754>.
- Silva, P. (2021). Possessing reasons: why the awareness-first approach is better than the knowledge-first approach. *Synthese*. <https://doi.org/10.1007/s11229-020-02916-5>.
- Silva, P. (2020). Knowledge-First Theories of Justification. *Internet Encyclopedia of Philosophy*. <<https://iep.utm.edu/kft-just/>>.
- Silva, P. (2019). Beliefless Knowing. *Pacific Philosophical Quarterly*, 100(3), 723-746.
- Silva, P. (2018a). Justified Group Belief is Evidentially Responsible Group Belief. *Episteme*, 16(3), 262-281.
- Silva, P. (2018b). Explaining enkratic asymmetries: knowledge-first style. *Philosophical Studies*, 175(11), 2907-2930.
- Silva, P. (2017). Knowing how to put knowledge first in the theory of justification. *Episteme*, 14(4), 393-412.
- Silva, P. (2015). Does Doxastic Justification Have a Basing Requirement? *Australasian Journal of Philosophy*, 93(2), 371-387.
- Silva, P. (2013a). Epistemically self-defeating arguments and skepticism about intuition. *Philosophical Studies*, 164(3), 579-589.
- Silva, P. (2013b). How to Be Conservative: A Partial Defense of Epistemic Conservatism. *Australasian Journal of Philosophy*, 91(3), 501-514.
- Smithies, D. (2012). The Normative Role of Knowledge. *Noûs*, 46(2), 265-288.
- Smithies, D. (2019). *The Epistemic Role of Consciousness*. Oxford: Oxford University Press.



- Sosa, E. (2015). *Judgment and Agency*. Oxford: Oxford University Press.
- Stanley, J. (2005). *Knowledge and Practical Interests*. Oxford: Oxford University Press.
- Stanley, J. (2011). *Know How*. Oxford: Oxford University Press.
- Stanley, J. and T. Williamson (2001). Knowing How. *Journal of Philosophy*, 98(8), 411-444.
- Sylvan, K. (2018). Knowledge as a Non-Normative Relation. *Philosophy and Phenomenological Research*, 97(1), 190-222.
- Turri, J. (2010). Does Perceiving Entail Knowing? *Theoria*, 76, 197-206.
- Turri, J. (2011). The Express Knowledge Account of Assertion. *Australasian Journal of Philosophy*, 89(1), 37-45.
- Turri, J. (2015). Unreliable knowledge. *Philosophy and Phenomenological Research* 90(3): 529–45.
- Unger, P. (1972). Propositional verbs and knowledge. *Journal of Philosophy*, 69(11), 301-12.
- Unger, P. (1975). *Ignorance: A Case for Scepticism*. Oxford: Clarendon Press.
- van Woudenberg, R. (2009). Ignorance and Force: Two Excusing Conditions for False Beliefs. *American Philosophical Quarterly* 46(4): 373-386.
- Vetter, Barbara. (2015). *Potentiality*. Oxford: Oxford University Press.
- Wallis, C. (2008). Consciousness, Context, and Know-How. *Synthese*, 160(1), 123-153.
- Williamson, T. (2000). *Knowledge and Its Limits*. Oxford: Oxford University Press.
- Zimmerman, M. (2008). *Living with Uncertainty: The Moral Significance of Ignorance*. Cambridge University Press.
- Zimmerman, M. (2020). Willful Ignorance and Moral Responsibility. In Mark Timmons, ed., *Oxford Studies in Normative Ethics*, Vol. 10. Oxford: Oxford University Press.