

AWARENESS BY DEGREE

Paul Silva Jr.
University of Cologne

Robert Weston Siscoe
Bowling Green State University

Abstract. Do factive mental states come in degrees? If so, what is their underlying structure, and what is their theoretical significance? Many have observed that ‘knows that’ is not a gradable verb and have taken this to be strong evidence that propositional knowledge does not come in degrees. This paper demonstrates that the adjective ‘aware that’ passes all the standard tests of gradability, and thus strongly motivates the idea that it refers to a factive mental state that comes in degrees. We will explore a range of structural questions that have been overlooked in regard to our awareness of facts, showing that the degreed structure of awareness has significant implications for our understanding of knowledge.¹

¹For helpful discussion and feedback on earlier drafts of this paper, we are grateful to Eli Alshanetsky, Philip Atkins, Peter Baumann, Marvin Backes, Sven Bernecker, Thomas Grundman, Luis Rosa, two anonymous referees, and audiences at the University of Cologne, the University of Glasgow, the University of Hamburg, the University of Luxembourg, Temple University, George Fox University, and the 2023 Central APA.

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.

David Chalmers (1996: 29)

1. Introduction

Natural languages contain a rich and varied inventory of factive terms. To start with, there are **factive stative verbs**. These verbs signal that one has come to be in a mental state that represents a fact (rather than a falsehood) as obtaining.² Core examples of factive stative verbs include: ‘knows that’, ‘recognizes that’, ‘learned that’, ‘discerned that’, ‘discovered that’, and many others. These are called *factives* because they refer to mental states that require the truth of their content. You cannot, for example, know *that Mont Blanc is large* if that is false.

In addition to factive stative verbs, natural languages contain **factive stative adjectives**. Unger (1975: 171-176) drew attention to a wide range of *emotive* factive adjectives (‘happy that’, ‘angry that’). But there are also many *epistemic* factive adjectives. Here are a few:

- ‘aware that’ / ‘aware of the fact that’
- ‘conscious that’ / ‘conscious of the fact that’
- ‘obvious that’
- ‘clear that’

There are many questions to ask about the factive adjectives. We will be addressing only a very specific set of these questions concerning factual awareness and its relation to knowledge in what follows. The set of questions raised here are of fundamental importance to epistemology as a

²Sometimes the term ‘fact’ is connected to truth-makers and taken to refer to some *particular, x, exemplifying some property F*, e.g. *the computer’s being large*. But in what follows we’re concerned with the sense of ‘fact’ connected to truths and involve *that-clauses*, e.g. being aware of (or ignorant of) *the fact that the computer is large* (cf. Frege 1977:25; Strawson 1998:403; cf. Hyman 2017). That-clause facts are expansive: they can involve facts involving quantifiers, negatives, disjunctions, and the like. In contrast, ‘facts’ as exemplifications of properties seem 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.

striking and underexplored asymmetry between the verb ‘knows that’ and the adjective ‘aware [of the fact] that’ concerns the question of **gradability**. Many theorists have observed that ‘knows that’ is not gradable. In contrast, consider **factual awareness**, i.e. the awareness relation referred to in English with the expressions ‘is aware that p ’ and ‘is aware of the fact that p ’. Graded expressions involving factual awareness are semantically unproblematic and examples are numerous and easily found in use among ordinary English speakers. Further indirect evidence of the gradability of awareness-encoding terms is found in the large literature on degrees of *non-factual* awareness in cognition, e.g. visual awareness of concrete objects (Fazekas & Overgaard, 2018). Unfortunately, issues surrounding the degreed structure of factual awareness are complex. So space prohibits a comprehensive discussion of degrees of non-factual awareness, the potential gradability of the other factive adjectives noted above, as well as the degreed character of *consciousness*, where the notion of consciousness here is intransitive (it takes no object), it’s non-factive, and it’s tied to phenomenal experience (Lee, 2022). But we believe that making clear and concrete progress on the wide range of issues associated with degrees of factual awareness and their relation to propositional knowledge will promote a more effective and efficient exploration of the gradability of other factive terms and the relations they refer to.

Here’s the roadmap. Section 2 examines the evidence for the non-gradability of ‘knows that’. Section 3 provides a comprehensive case for the gradability of ‘aware [of the fact] that’ and provides an account of the grounds of degrees of awareness. Section 4 explains four ways of modeling degrees of factual awareness from linguistics. Section 5 provides a set of considerations that favor one of the four models. Section 6 explores the relation between knowledge and degrees of factual awareness.

2. Propositional Knowledge Does Not Come in Degrees

Does (propositional) knowledge come in degrees? This question is much more significant than it might initially seem. First, if knowledge doesn’t come in degrees but factual awareness does, then we cannot identify the two factive relations. Second, if ‘knows that’ is not gradable, but intimately connected to a degreed factive relation (like factual awareness), then new options come into view for understanding the nature of knowledge. This will be picked up in Section 6. But first things first. This section will explain why the balance of evidence suggests that propositional knowledge does not come in degrees. The next section shows that there is a powerful case to be made that factual awareness comes in degrees.

So, are there degrees of knowledge? Gilbert Ryle (1949:59) asserted that ‘we never speak of a person having partial knowledge of a fact or truth [...] either he knows this fact or he does not know it’, and Fred Dretske (1981: 363) held that ‘knowing that something is so, unlike being wealthy or reasonable, is not a matter of degree’. This is by no means a minority position, and for good reason.³ In a systematic defense of the non-gradability of ‘knows that’, Jason Stanley (2004, 2005: Ch.2) points out that ‘knows that’ lacks almost all of the characteristic features of semantic gradability. The first marker mentioned by Stanley is that gradable terms accept degree modifiers, terms like ‘somewhat’, ‘fairly’, ‘very’, ‘quite’, and ‘rather’. Take the gradable adjective ‘tall’ as an example. A person can be somewhat, fairly, or very tall:

(1) Lucy is somewhat [/fairly/very] tall.

Along with accepting degree modification, gradable terms like ‘tall’ can also be used in comparative constructions:

(2) Lucy is taller than John.

According to Stanley, ‘knows that’ does not exhibit either of these markers of gradability. Now some have noted that ‘tall’ is an adjective rather than a verb and that ‘tall’ doesn’t take the sentential complement ‘that *p*’, so there are multiple disanalogies between ‘tall’ and ‘knows’. But it is equally important to note, as Stanley does, that ‘knows’ differs even from other verbs that accept sentential complements, including verbs like ‘regrets that’, ‘suspects that’, and ‘hopes that’.

Notably, the features of gradability exhibited by ‘tall’ are also exhibited by these verbs that take sentential complements, since they also accept degree modifiers and can feature in comparative constructions:

(3) Lucy slightly [/very much/strongly] regrets that she did not graduate from high school.

(4) Lucy regrets that she didn’t graduate from high school more than she regrets that she didn’t graduate from college.

³Others who hold that propositional knowledge does not come in degrees include Bengson and Moffett (2011), Blome-Tillmann (2017) and (2022), Brogaard (2016), Crane (2012), Kvanvig (2003), Pavese (2017), Roberts (2009), and Schaffer and Szabó (2014). For minority positions on which knowledge does come in degrees, see Hetherington (2001, 2011) and Sosa (2001,2011).

We are now in a position to see that ‘knows that’ does not display the typical gradability that we find with other verbs and adjectives. To begin with, ‘knows that’ accepts *none* of the degree modifiers that we have mentioned, neither those that we consistently find with gradable adjectives nor those that frequently pair with gradable verbs:

(5) #Lucy slightly [/somewhat/fairly/very/highly/strongly] knows that smoking causes cancer.

Furthermore, knowledge does not seem to allow for comparative constructions:

(6) #Lucy knows that it’s raining more than she knows that it’s hailing.

Given that both (5) and (6) are infelicitous, it appears that Stanley is right to think that ‘knows that’ does not exhibit the two key markers of gradability.

Not everyone has accepted Stanley’s argument that ‘knows that’ does not display any signs of gradability. In particular, some authors have pointed out that ‘knows that’ seems to accept degree modification by variations of ‘well’:⁴

(7) Lucy knows well [very well/quite well/full well] that smoking causes cancer.

With (7), we can see that ‘knows that’ accepts modification by terms and phrases that, at first glance, appear to function with a degree scale of some sort, as it is possible to not only know that *p*, but also to know very well that *p* and know full well that *p*. Linguists have also observed that modification by ‘very’ and ‘well’ is typically complementary. If a term combines with ‘very’, it typically does not combine with ‘well’ and vice versa.⁵ Thus, it might be that ‘knows that’ is gradable after all, albeit without accepting the characteristic degree modifiers that Stanley originally had in mind.

There are two plausible ways to interpret the ‘well’-modification of knowledge, but neither of those options implies that knowledge itself comes in degrees. The first option is that modifying ‘knows that’ with ‘well’ introduces a new sort of degree scale. Take, for example, the deverbal adjective ‘written’ and the modified adjectival phrase ‘well written’. With ‘written’ there is a scale

⁴ See Dutant (2007), Hetherington (2001) and (2011), Lai (2021), Pavese (2017), Schaffer and Szabo (2014), and Stanley (2004) and (2005: Ch.2).

⁵ See Kennedy and McNally (2005) and Knowles (1974). For example, ‘He is very tall’ is acceptable while ‘He is well tall’ is not, and ‘The case was well documented’ is more natural than ‘The case was very documented’. This does not preclude the possibility that, once ‘well’-modification introduces another scale, ‘very well’ becomes acceptable, as with ‘The case was very well documented.’

associated with whether a text is partially, almost, or fully written. This scale concerns the extent to which the text is complete. The scale introduced by ‘well’, on the other hand, is about *the quality* of the written text. It’s a different thing altogether to talk about an article that is completely written and an article that is well written.

This suggests that ‘well’ introduces a new sort of scale. An interesting fact about the modifier ‘well’ is that it allows for further degree modification:⁶

(8) The article was somewhat [/fairly/very] well written.

(9) The cargo truck was somewhat [/fairly/very] well loaded.

The same is not true of other degree modifiers like ‘somewhat’:

(10) #Lucy is slightly [/somewhat/fairly/very] somewhat tall.

(11) #The road is slightly [/somewhat/fairly/very] somewhat long.

The fact that ‘well’ easily accepts modification by ‘slightly’, ‘fairly’, and ‘very’ suggests that felicitous uses of ‘well’ introduce a new scale (and hence genuine gradability), and that this new scale can be introduced *whether or not* the target term was gradable and hence associated with another scale.

Because ‘well’ makes use of a scale that is orthogonal to the original degree scale (if there was one), Julien Dutant (2007) has distinguished between ‘quantitative gradability’, the sort of gradability possible for the original, unmodified term, and ‘qualitative gradability’, the degree scale of quality introduced by ‘well’. While *qualitative* gradability can be used to offer assessments of the quality of something, like the quality of a piece of writing, *quantitative* gradability refers to whether it’s actually possible to have differing amounts or degrees of that thing, like when a text is only partially written. Dutant then uses this distinction to argue that, while ‘knows that’ may admit of qualitative gradability, it ultimately does not allow for quantitative gradability. For example, suppose Lucy knew that a tree is behind the cafe on the basis of reliable testimonial evidence. Even so, she might get even better evidence for thinking that at some later time, e.g. she might go out and look at and touch the tree itself, and thereby come to have both testimonial and perceptual evidence – each individually sufficient for knowledge. *What has changed is not Lucy’s degree of knowledge, but the quality of her grounds for knowledge.* Dutant (2007: 13) appeals to such considerations to explain the qualitative gradability of ‘knows that’ associated with ‘well’ modifiers.⁷

⁶See Kennedy and McNally (2005) and (2012).

⁷Lai (2021:3954) has resisted Dutant’s idea that ‘well’ encodes a separate scale of quality, objecting that paradigmatically degreed notions like ‘understanding’ only accept qualitative degree modifiers and that sentences like ‘John understands a

Along with making sense of the limited gradability of ‘knows that’, Dutant’s distinction between typical quantitative scales and the scale introduced by ‘well’ allows us to accommodate accounts on which knowledge can be of better or worse quality without also conceding that knowledge comes in degrees. On Stephen Hetherington’s (2011: 172) view, for example, knowledge can come in degrees because we can know to a greater or lesser extent how it is that p . But as George Pappas (2017: 310) observes, we can still take on this thought that a ‘person’s knowledge of the Pythagorean theorem might be firmer than that of another, in that she knows how to prove that theorem within Euclidean geometry’ without also saying that one person has a greater degree of knowledge. Instead, we can follow Dutant’s idea: even though ‘knows that’ is not quantitatively gradable, the quality of knowledge can nevertheless be improved.

Dutant’s approach, however, is not the only way to interpret the ‘well’-modification of ‘knows that’. We could follow Stanley (2004: 126, 2005: 38-39) in holding that phrases like ‘knows well’ and ‘knows very well’ could simply be *used emphatically*. One reason to think that this is what is happening in examples like (7) is that, as emphasized by Jonathan Schaffer and Zoltán Szabo (2014), the quality of propositional knowledge is not adjustable all the way up and down the scale. To begin with, though we can say that a text is somewhat or partially well-written, the same does not seem true with ‘knows well that’:

(12) The article was somewhat [/partially] well written.

(13) #Lucy knows somewhat [/partially] well that smoking causes cancer.

Furthermore, with the quality scale of writing, it’s also possible to introduce a scale of quality using the negatively charged ‘poorly’ and ‘badly’:

(14) The article was poorly [/badly] written.

(15) The cargo truck was poorly [/badly] loaded.

As with ‘well’, ‘poorly’ and ‘badly’ don’t select for degrees on the quantitative scale. Instead, they introduce a new scale of quality, allowing further degree modification on that quality scale by terms like ‘somewhat’, ‘fairly’, and ‘very’. We cannot, however, pair ‘knows that’ with ‘poorly’ or ‘badly’:

lot that p ’ are infelicitous. However, Lai’s argument is questionable. For in considering degrees of understanding, Hills (2016:665-666) points out that ‘understands why’ accepts both qualitative degree modifiers (John understands fairly/very well why smoking causes cancer) and quantitative degree modifiers (John barely/partially/fully/completely understands why smoking causes cancer).

(16) #Lucy poorly [/badly] knows that smoking causes cancer.

Further supporting the idea that ‘well’-modification of ‘knows that’ is merely pragmatic is that it sounds borderline contradictory to attempt to contrast knowing and knowing well. With both ‘written’ and ‘loaded’, we can distinguish between the job being merely complete and it being completed well:

(17) The article is written, but it’s not written *well*.

(18) The cargo truck is loaded, but it’s not loaded *well*.

The same is difficult, if not impossible, to do with knowledge:

(19) #Lucy knows that smoking causes cancer, but she doesn’t *know it well* [/know it *full well*].⁸

All of this suggests that, when it comes to knowledge, ‘well’-modification might only be usable for pragmatic emphasis. Needless to say, if uses of ‘knows well’ and ‘knows full well’ are merely for pragmatic emphasis, then such constructions don’t ensure that propositional knowledge comes in degrees.⁹ Thus, we have two possible ways of interpreting ‘well’-modification of ‘knows that’, both in favor of the conclusion that knowledge does not come in degrees.

3. Factual Awareness Comes in Degrees

Are there degrees of factual awareness? The evidence that factual awareness comes in degrees is of the same kind and strength as the evidence that other epistemic notions can come in degrees.¹⁰ But before providing that evidence several clarifications will aid the discussion to follow.

⁸Sosa (2011) introduces a set of conditions for knowing which, when satisfied by an agent, the agent is said to not only know, but to ‘know full well’. This is not a problem for us. Obviously, Sosa’s *stipulative* use of the expression ‘knows full well’ allows Sosa to meaningfully assert that ‘S knows that *p*, but she does not know *p* full well’. Our concern is not with stipulative uses of this expression.

⁹What this pragmatic emphasis accomplishes may differ depending on the case. If someone says that they know very well that smoking causes cancer, this may be to emphatically respond to a challenge to their knowledge of the consequences of smoking. If, on the other hand, someone says ‘Lucy knows better than most that smoking causes cancer - she has had cancer herself!’, they may convey that Lucy’s evidence is more salient or immediately available than others that also know that smoking causes cancer.

¹⁰See, for example, work on degrees of rationality (Siscoe, 2021b, 2022a, 2022b, and 2023; Sorensen, 1991), degrees of justification (Hawthorne & Logins, 2020; Fassio & Logins, 2023; Siscoe, 2021a), degrees of certainty (Lassiter, 2017; Beddor, 2020), and degrees of ignorance (Brogaard, 2016; Siscoe, 2024).

First, we will continue to use the braced expression ‘aware [of the fact] that’ for convenience to avoid separately writing out ‘aware that’ and ‘aware of the fact that’. Like others, we take there to be no semantic difference between these expressions.¹¹ To say ‘She is aware that he arrived’ is to say no more and no less than that ‘She is aware of the fact that he arrived.’

Second, despite our focus on factual awareness it is worth pointing out that there are various objects the adjective ‘aware of’ can take.¹² You can be aware of *particulars* (e.g. aware of *the kangaroo*), you can be aware of *properties* of particulars (e.g. aware of *the smell* of the kangaroo), you can be aware of *events* (e.g. aware of *the departure* of a kangaroo), you can be aware of *propositions* (e.g. aware of *the meaning of the sentence* ‘some kangaroos live in kitchens’), you can be aware of *facts* (e.g. aware [of *the fact*] that kangaroos are non-human animals), you can be aware of *qualia* (e.g. aware of *what it’s like* to feel a kangaroo’s fur), and you can be aware of *skills* (aware of *how to scare kangaroos*). In this way, ‘aware of’ is like ‘knows’, which also takes various objects. You can know *persons*, know *how* to do something, know *the meaning of the sentence* ‘Elia enjoys Susan’, know *what it’s like* to do something, and know *that p*. In what follows our sole concern is with the relation between propositional knowledge and factual awareness.

Thirdly, and as others have observed, *awareness of facts* should not be conflated with *consciousness of facts* (Chalmers, 1996: 28-29, 221-222; Silva, 2023, Forth). Being conscious of a fact *p* implies conscious experience in some way. But the awareness of facts does not. As Chalmers (1996: 221-22) points out: ‘I may be (nonoccurrently) aware that there is a bicycle downstairs, without there being an associated bicycle experience. This sort of awareness without experience is most pronounced with propositional [=factual] awareness...’ Factual memory serves as an apt example. You learned, memorized, and have not since forgotten your multiplication tables. So you are aware of the fact that the product of 6 and 4 is 24. But you were aware of that fact even *before* this very example made you conscious of it by directing your attention towards it. Similarly, throughout reading this paper you will be aware of the fact *that this is a philosophical text with more than ten footnotes*.¹³ But this fact will surely drift out of consciousness not long after reading this section. In this way being aware of the fact that *p* should be distinguished both from *being conscious of* the fact that *p* as well as *attending to* the fact that *p*. Further, since knowing that *p* entails being aware of the fact that *p*, factual awareness cannot be more demanding than knowledge.¹⁴ Thus, since it’s possible to know

¹¹Dretske (1993), Littlejohn (2015:598), Silva (2023), and Silva and Siscoe (2024).

¹²Many have observed this, including Dretske (1993), Chalmers (1996:28ff,221ff), Huemer (2001), Littlejohn (2015), Silva (2023), Silva and Siscoe (2024).

¹³Thanks to Marvin Backes for pointing this out.

¹⁴ Many epistemologists have explicitly held that if *S* knows that *p*, then *S* is aware [of the fact] that *p*. See Dretske (1993), Huemer (2001), Littlejohn (2015), Jennifer Nagel (2017), and Silva (2023, Forthcoming). For more on the knowledge-awareness relation see also Section 6 of this paper.

that p while asleep and unconscious, it follows that one can be aware of the fact that p while asleep and unconscious. The same is not true of being conscious of the fact that p . Finally, cognitive scientists and philosophers study *unconscious perception* (Prinz, 2010; Block, 2016). Since the concept of perception is a concept of one *way* of being aware – i.e. perceptual awareness – the study of unconscious perception is the study of unconscious awareness. While the object of their studies is typically perceptual awareness of particulars, this still provides evidence that the general concept of awareness is to be disassociated from the concept of consciousness.

Fourth, questions about degrees of factual awareness should be distinguished from questions about ‘awareness growth’. The literature on awareness growth has to do with one’s awareness of *propositions* of which one was previously unaware (Steele & Stefánsson, 2021; Canson, 2024). For example, before you had the concepts PRIME NUMBER and FAVORITE you were unable to understand the meaning of the proposition: *everyone has a favorite prime number*. But once you obtained those concepts and became aware of *the meaning* of the statement (a proposition), your ‘awareness grew’ in the sense that you became newly aware of a proposition. But this kind of awareness – propositional awareness – is non-factive, and no one working in the awareness growth literature maintains that the awareness of propositions is factive. So the awareness of propositions (i.e. the meanings of sentences) does not imply awareness of facts, even if the proposition one is aware of happens to be true.

In what follows we will provide a wide range of examples where factual awareness is taken to come in degrees and where both ‘aware that’ and ‘aware of the fact that’ are used in comparative forms. After that, we will discuss *the grounds* for degrees of awareness, i.e. what often explains the truth of degreed awareness claims.

Let’s start with quantitative degree modifiers. Observe that all of the following are felicitous:

(20) Lucy is slightly [/somewhat/very/highly/fully/completely] aware [of the fact] that it is raining.

Examples can be found in recent popular news sources:

- ‘I’ve known him for years and was only *slightly aware* that he wrote and directed musicals.’
-[San Francisco Chronicle](#), November 2016
- ‘He had already been *somewhat aware* that he was at risk for heart disease after two older brothers had died of heart attacks, one in his 40s.’
-[U.S. News and World Report](#), November 2021

- ‘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.’

-[Reuters](#), November 2021

- ‘Countless other black students who tried to take advantage of integration in the ‘60s were *highly aware* that these predominantly white schools were physically unsafe spaces for them.’

-[The Atlantic](#), November 2015

- ‘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.’

-[The Wall Street Journal](#), August 2021

Along with accepting this wide variety of degree modifiers, factual awareness can also feature in comparatives:

(21) Lucy is more aware [of the fact] that smoking causes cancer than she is that sunbathing causes cancer.

(22) Every day I am becoming increasingly more aware [of the fact] that smoking causes difficulty breathing.

In addition to these sorts of comparatives, stronger and weaker forms of awareness may be contrasted. Being slightly aware or somewhat aware is clearly distinct from being fully aware:

(23) Lucy is somewhat aware [of the fact] that sunbathing causes cancer, but she’s not fully aware [of the fact] that it does.

(24) John is slightly aware [of the fact] that it is raining, but he is not completely aware [of the fact] that it is.

Clearly, ‘aware [of the fact] that’ displays two important markers of semantic gradability emphasized by Stanley: it accepts degree modifiers and can be used in comparative constructions.

While these features provide strong evidence that factual awareness comes in degrees,¹⁵ we will provide a series of cases to help us better understand what degreed properties might underlie and explain the correctness of quantitative and comparative statements about factual awareness. Consider the following:

Knowing Noel

Noel is walking by the small, cramped supply closet next to Bart's office and sees the janitor looking for something. In the process of searching, the janitor bumps into the wall with their shoulder, making a loud noise. Seeing all this, Noel comes to believe and to know *that a person caused the noise by bumping into the wall*.

Seeing Susan

Like Noel, Susan is walking by the small, cramped supply closet next to Bart's office and sees the janitor looking for something. In the process of searching, the janitor bumps into the wall with their shoulder, making a loud noise. Susan sees all of this happen, and for this reason it *seems to Susan that* that the noise was caused by a person bumping into the wall. But Susan *does not come to believe this* (and thus does not come to know it). For she just left the psychology department's experimental lab where she was given a pill prior to performing a series of audio-visual identification tasks. Prior to taking the pill she signed a waiver that indicated the following: all participants would have a 65% chance of being given a drug that would seamlessly distort one's visual and auditory experiences in undetectable ways. But there is a 35% chance of being given a placebo. But the whole study was a hoax, no one was given any

¹⁵It's also worth noting that not only does factual awareness seem to display quantitative gradability, it demonstrates qualitative gradability as well. Like 'knows that', factual awareness can also be modified by 'well': (i) Lucy is well/very well aware [of the fact] that smoking causes cancer. While it's helpful to acknowledge that factual awareness also accepts qualitative gradability, we mention it only to set it aside. Like with propositional knowledge, there is a serious concern that 'well' modification of factual awareness is merely emphatic. To begin with, it's not clear that this sort of evaluation exists up and down a qualitative scale. While 'poorly aware that/of the fact that' sounds acceptable, saying that someone is 'somewhat well aware' or 'partially well aware' is forced at best: (ii) #Lucy is somewhat/partially well aware [of the fact] that smoking causes cancer. Furthermore, it seems doubtful that we can contrast awareness simpliciter with well awareness: (iii) #John is aware [of the fact] that the dog is barking, but he is not *well* aware [of the fact] that it's. Despite these considerations, it could be that the qualitative gradability of factual awareness is genuine. Nevertheless, because the 'well' modification of factual awareness might be merely pragmatic, it will play no role in our case that factual awareness comes in degrees.

perception-distorting drug. While Susan doesn't know that it was a hoax, she *knows that cognitive-psychology experts indicated that she was much more likely than not to get a distortion-inducing pill.*

Bothered Bart

Bart is sitting in his office and hears a single, solid bump against the wall. Bart knows that there is a small, cramped supply closet on the other side of the wall, and he knows that janitors are regularly in there to get cleaning supplies. Bart knows that *the bump-sound might have been caused by a person bumping into the wall* and he knows that *it might have been caused by a falling stack of cleaning supplies that hit the wall.* Bart knows that both kinds of events have frequently caused bumping noises in the past. But Bart does not have any evidence as to which of the potential causes is responsible for the noise that he heard. The fact of the matter is that the noise was caused by a janitor bumping into the wall.

Clueless Carl

Carl lives in a rural village on the other side of the world with no knowledge whatsoever of Bart, Bart's office, the noise Bart just heard, nor the fact that there is a supply closet adjacent to his office.

Take the fact *that the noise was caused by a person bumping into the wall.* Noel, Susan, Bart, and Carl stand in some epistemic relation to this fact, but how should we understand this difference in epistemic position? There are a few judgments that are natural to make about these cases.

A1 Clueless Carl is *completely unaware* of this fact. That is, he is *not at all aware* [of the fact] that the noise was caused by a person bumping into the wall.

Knowing Noel, on the other hand, saw the whole event in a perfectly normal perceptual environment and with perceptual capacities that were functioning normally. Further, Noel (unlike Susan) had no reason to think that she was (or might be) in a situation where her perceptual capacities were liable to mislead her. For these reasons, it's natural to think that Noel *knows that the noise was caused by a person bumping into the wall.* It's also natural to think that:

A2 Noel is *completely aware* [of the fact] that the noise was caused by a person bumping into the wall.

We will explore some reasons in support of connecting knowledge and complete awareness in Section 6. For now we just want to register the fact that it is intuitive to say that Noel is completely aware of the fact that the noise was caused by a person bumping into the wall.

Now contrast Clueless Carl with Bothered Bart and Knowing Noel. Carl is not at all aware of the fact *that the noise was caused by a person bumping into the wall*, while Noel is completely aware of it. But Bart is somewhere in between. In particular:

A3 Bart is *not completely unaware* [of the fact] that the noise was caused by a person bumping into the wall. Rather, Bart is *slightly* or *somewhat aware* of that fact.

Denying A3 is difficult. Bart is clearly better placed than Carl in regard to that fact. Bart (i) *heard the noise* and (ii) *Bart knows that the noise might have been caused by a person bumping the wall* because he *knows that this is a very frequent and normal cause of this kind of sound*.¹⁶ While we don't want to say that Bart is aware of that fact *simpliciter* (something we will defend in Section 5), we do want to say that Bart is slightly or somewhat aware of that fact. To deny this is to invite the idea that Bart and Carl stand in the same awareness relation to that fact, despite (i) and (ii). This strikes us as very counterintuitive.¹⁷

Now contrast Bart and Susan. Both are in a *worse* epistemic position than *Knowing Noel* in relation to the same fact. But Susan is also in a better position than Bart. After all, Susan saw and heard the janitor bump into the wall, and she was in perfectly normal perceptual circumstances. She just had the bad luck of having a (misleading) reason to think her perceptual capacities were impaired by a drug. So we judge that:

¹⁶ A referee wondered if something like the following was true: *Necessarily, if p is in fact true and S is aware of the fact that p might be true, then S is to at least some degree aware of the fact that p .* Our major conclusions in this paper concerning degrees of awareness are consistent with the affirmation and denial of this claim. Against this, many have thought that factual awareness and knowledge are alike in having some kind of non-accidentality constraint, and so the plausibility of this principle, as far as we can tell, will depend on whether or not it is always the case that: those things, m , that make S aware that p might be true are suitably non-accidentally related to the fact that p . It seems possible to construct cases where m – i.e. that which makes one aware that p might be the case – is only accidentally related to the fact that p even though m is non-accidentally related to the fact that p might be the case.

¹⁷ A referee helpfully pointed out that the following sounds odd: Bart is *somewhat aware* that a person bumped the wall, but Bart is not aware (*simpliciter*) that a person bumped the wall. But this claim is ambiguous. One can be somewhat aware while having a comparatively weak degree of awareness *or* somewhat aware while having a significantly higher degree of awareness (so long as it's not too close to the maximum degree of awareness). As we will discuss below, partial absolute gradable adjectives, F , that have liminal regions allow for low degrees of F sufficient for being *somewhat F* without being *F simpliciter*.

A4 Susan is *somewhat more aware* than Bart [of the fact] that the noise was caused by a person bumping into the wall.

One referee suggested, in contrast, that Bart might be more aware of the Bumping Fact than Susan because Bart ‘can rule out’ more possibilities than Susan. But this is far from obvious if our concern is the extent to which their evidence ‘rules out’ various possibilities. Consider the following relation:

Evidential Ruling Out. For any proposition p , p is evidentially ruled out by the facts, E , that one has access to iff: p is incompatible with E .

This is at least one perfectly respectable and reasonably familiar notion of evidence ruling out a hypothesis. Because Susan is in normal visual circumstances she *visually perceives the janitor bumping the wall* and also *auditorily perceives the resulting noise* (object perception). In addition to this, many think that *perceiving that p* (factual perception) is compatible with lacking sufficient reason to believe that p ; specifically, many think that misleading higher-order defeaters of the sort Susan has don’t prevent her from *perceiving that* the bumping noise was caused by a person.¹⁸ Rather, these defeaters at most prevent her from having *sufficient* reason to believe that the bumping noise was caused by a person.¹⁹ For this reason, it’s defensible to hold that Susan’s seeing that the bumping noise was caused by a person, gave her access to that fact and put it into her ‘evidence bank’. So, arguably, Susan’s evidence includes the fact that a person caused the noise, and possessing that fact evidentially rules out all propositions incompatible with it. In contrast, Bart’s evidence does not include this fact, nor does his evidence rule out all propositions incompatible with it. This is because Bart only *auditorily perceives the noise*, and *knows that the janitor is one of (at least) two possible causes*. Neither puts Bart in possession of facts that evidentially rule out every proposition incompatible with the fact that a person caused the noise.

For expressive convenience, we’ll start using the expression **the Bumping Fact** in place of ‘the fact that the noise was caused by a person bumping into the wall’. Putting A1-A4 into a picture, we have the following ordering of degreed awareness relations involving the Bumping Fact:

¹⁸ For endorsements and defense see McDowell (2002:277), Pritchard (2012b:26), Turri (2010), Schroeder (2021:Ch.4), and Silva (2023:51-54;Ch.6); cf. Bernecker (2010:78).

¹⁹ A noteworthy consequence of this is that one can possess the fact that p as evidence while lacking sufficient reason to believe p . But this is a familiar lesson from the literature on self-defeating propositions, i.e. propositions one has conclusive evidence for (e.g. propositions entailed by what one knows) *only so long as one fails to believe them*. See Silva and Bernecker (2023) for discussion and references on the distinction between *having evidence that supports p* and *having sufficient reason to believe p* .




Figure 1. An Awareness Ordering

It's unclear to us whether the associated scale of awareness has a maximum degree, and we use *the black diamond* to indicate this.

What properties ground this ordering? And, more generally, what grounds degrees of awareness? Reflection on the cases above from both news sources and our constructed cases with the Bumping Fact suggest two related properties essential to understanding what grounds the extent to which a subject is aware of the Bumping Fact. The first has to do with whether or not and to what extent each subject **possesses reasons for believing** the Bumping Fact. Clueless Carl has no reason whatsoever to believe the Bumping Fact. And each person along the line of Figure 1 has increasingly strong reasons that support believing the Bumping Fact. Bart has *some*, though not much, reason to believe the Bumping Fact. Susan has *much more* reason to believe it. While Noel has *even more* reason to believe it than Susan.²⁰

But facts about degrees of factual awareness are not solely tied to increasingly strong reasons for belief. For in order for one's reasons to relate one to a fact **one's reasons must be suitably non-accidentally related to that fact** (cf. Littlejohn, 2017: 26-27). For example, Bart's reasons for believing the Bumping Fact are non-accidentally related to it. For it is *because* the Bumping Fact obtains that Bart has *some reason* to believe the Bumping Fact. Similarly for Susan: it is *because* the Bumping Fact obtains that Susan has *significant reason* to believe the Bumping Fact. Similarly for Noel: it is *because* the Bumping Fact obtains that Susan has *significant reason* to believe it, and it is because she has no further countervailing reasons that she has *sufficient reason* to believe it.

Notice that having non-accidental reasons for belief is an implicit part of the backstory in the sentences of degreed awareness noted above:

- 'I've known him for years and was only *slightly aware* that he [Morris Bobrow] wrote and directed musicals.'

²⁰ One referee wondered if both Susan and Bart could have *no reason* to believe the Bumping Fact. While neither may have sufficient reason to believe the Bumping Fact, they certainly have *more reason* to believe the Bumping Fact than they have to believe obvious falsehoods and facts they have no access to whatsoever. So both must have *at least some* reason to believe the Bumping Fact, and Susan must have more reason than Bart for reasons noted above.

It would not only be very odd to assert this and then go on to assert that one had *no reason* to think Morris wrote or directed musicals. Furthermore, it would also be very odd to assert this and then go on to say that one's reasons were all the product of hallucination, wishful thinking, or some other source quite accidentally related to the fact that he wrote and directed musicals. Just suppose you asked the author, Beth Spotswood, how she came to be slightly aware of the fact that Morris wrote and directed musicals. If Beth responded by saying 'I once had a dream of Morris writing and directing musicals', that would be strange and not at all what we expected her to say. Rather, we would have expected her to produce weak evidence that is appropriately related to the fact that Morris wrote and directed musicals. For example, that Beth once saw 'Directing *Singing in the Rain*' on Morris' calendar, or that Beth once heard Morris assert 'Writing and directing musicals is arduous work', or some other sliver of evidence connected to the fact that Morris wrote and directed musicals.

Indeed, in every case where a degree of factual awareness is asserted we have found it hard to generate a coherent understanding of the case where (i) one intuitively *retains* some (any) degree of awareness, even though (ii) the fact that p is *only accidentally related* to the fact one has some [/significant/sufficient/decisive] reason to believe that p . There are, of course, very difficult questions about what relations count as 'non-accidental' for the purposes of degrees of awareness. But this is a general problem facing all theories of factive stative relations, and one of the most famous problems facing those attempting to give a theory of knowledge. So we will not attempt to solve this issue here.

Furthermore, we will not pursue the conjecture that all degreed awareness states are grounded in non-degreed factive states, e.g. states of factual awareness *simpliciter* (Silva, 2023), or knowledge states (Littlejohn, 2017), or states of epistemic certainty (Beddor, 2020). It is compatible with what follows that such non-degreed factive states alone provide one with a set of reasons R for believing p that can be non-accidentally related to p in the way required to generate degreed awareness judgments of the sorts observed above. Our aim is not to defend a complete theory of what grounds degrees of awareness, but to explore factors that clearly contribute to its degrees as well as the implications that the degreed structure of awareness has for theories of knowledge.²¹

²¹ Many gradable adjectives are multidimensional. While we treat 'aware that' as unidimensional and rely only on cases that can be ordered in a linear fashion, we don't rule out the possibility of multidimensionality. For that issue does not impact whether we should think of awareness as having a liminal and a limited region on its underlying scale, which is a primary concern in this paper.

4. Modeling Absolute Gradable Adjectives

It should now be clear that factual awareness comes in degrees. We can gain further insight into degrees of awareness by looking to the semantics literature to help us understand *what kind* of gradable adjective ‘aware’ is. There we find a well-studied and widely employed framework for analyzing gradable adjectives. The central idea is that every gradable adjective has an associated *scale*. For ‘tall’, the associated scale is *degrees of height*. For ‘expensive’, the associated scale is *units of cost*. For ‘dangerous’, the associated scale is *degrees of danger*.

Depending on how they interact with these underlying degree scales, gradable adjectives typically come in three different forms: **relative gradable adjectives** (RAs), **total absolute gradable adjectives** (TAs), and **partial absolute gradable adjectives** (PAs).²² Paradigmatic RAs include ‘tall’, ‘large’, ‘expensive’, and ‘long’. The denotations of RAs are determined using a contextually-determined threshold *that is typically fixed by the relevant comparison class*. For example, what counts as a *tall* athlete can depend on whether we are comparing people to the class of professional athletes, or the class of basketball players, or just athletes in general. The scales of RAs are also typically open-ended. With ‘tall’, it is always conceptually possible to be a few millimeters taller. So for the RA ‘tall’, we can provide a geometrical model as follows:




Figure 2. Relative Gradable Adjectives

In contrast to RAs, we have TAs that include terms like ‘perfect’, ‘pure’, and ‘complete’. TAs are unlike RAs in that their denotations don’t rely on a contextually-determined threshold that is set using a comparison class, but are instead associated with the endpoint of the underlying degree scale. For example, the TA ‘perfect (*simpliciter*)’ refers to the maximum conceptual degree of perfection. TAs have different kinds of associated antonyms, including PAs, with representative examples of PAs including ‘imperfect’, ‘impure’, and ‘incomplete’. There are a few different ways to geometrically model TAs and their corresponding PAs. According to the **Simple Model**, the TA picks out the uppermost point of the top-closed scale, while its associated PA refers *to the rest of the degree scale*. Here, we can see the Simple Model as applied to ‘perfect’ and ‘imperfect’:

²²For linguists who discuss these types of gradable adjectives, see Burnett (2014) and (2017), Kennedy (2007), Kennedy and McNally (2005), Rotstein and Winter (2004), and Rusiecki (1985).




Figure 3. Simple Model for Perfect/Imperfect

There are a few things to keep in mind when reading these diagrams. First, *the hollow dot* on the left marks *the zero-degree point* on the scale. Second, *the black arrow* at the end represents the fact that the scale associated with the adjective *lacks a maximum degree*, i.e. for some things it is possible for them to become increasingly imperfect without end.

Third, the brace indicates that if something is imperfect to *any* (relevant) *degree*, it is *imperfect (simpliciter)*. This is true of all TA/PA pairs that satisfy the Simple Model. Nevertheless, an important thing to keep in mind is that some degrees of imperfection are *irrelevant* in some contexts. For instance, although full perfection only occurs at the endpoint we can still sometimes use ‘perfect’ to refer to things that have some slight imperfections. This is because TAs can be used in connection with *different standards of precision*.²³ For example, suppose an informal family meeting is supposed to begin at 5pm. We can count someone as having arrived on-time even if they arrive 20 seconds after 5pm since that is good enough for the practical purposes of an informal family meeting. But if we’re launching spacecraft and you’re essential to the launch, we’ll count you as being late if you’re not sitting at your station and ready to go at precisely 5pm. Similarly, we can often say that something is perfect so long as its slight imperfections are irrelevant. So while the zero-point on the scale excludes every *relevant* degree of imperfection, it can include irrelevant degrees of imperfection.

Even though many TA/PA pairs can be represented using the Simple Model, some pairs are more complex. Take the TA/PA pair ‘dry’ and ‘wet’. If we are in a situation where we must be exceptionally precise with our use of those terms, then we may be in a context where ‘dry’ and ‘wet’ seem to fit the Simple Model. For example, when it comes to food analysis, the dry matter of a portion of food is the food’s mass when free of all water. So for a portion of food to consist solely of dry matter it can contain no water whatsoever. In contexts where food analysts are concerned with dry matter, the Simple Model might adequately capture the kinds of judgments food analysts

²³For work on imprecision generally, see Krifka (2002) and (2007), Lakoff (1973), Sadock (1977), and Sauerland and Stateva (2007). For imprecision as it relates to absolute gradable adjectives, see Burnett (2014) and (2017), Kennedy (2007), Kennedy and McNally (2005), Lewis (1979), and Pinkal (1995).

want to make about portions of food when they are applying the most exacting standards for dryness/wetness.

Nevertheless, in most contexts, the typical ways that we apply ‘wet’ and ‘dry’ fail to accord with the Simple Model. Consider, for instance, a mildly damp towel. If we were using the food analysts’ standards for wetness, then the towel would count as wet. But as Rotstein and Winter (2004:265) observe, ‘the pair wet and dry is not necessarily complementary in all contexts. For instance, in some contexts a moist towel may be deemed neither wet nor dry.’ Or consider the TA ‘naked’ and its antonym ‘dressed’. A person is naked if they are not wearing any clothes. So wearing a bathing suit will prevent someone from being naked—at least according to *our* usual standard of precision. At the same time though, in many cases it would be incorrect to call someone ‘dressed’ if they are only wearing a speedo.

Rotstein and Winter’s point is that the Simple Model fails to accurately represent the wide range of cases where TA/PA pairs allow for a space between their extensions. Even though TAs like ‘naked’ and ‘dry’ refer to the end point on the associated degree scale, PAs like ‘dressed’ and ‘wet’ don’t seem to refer to *the entire* remaining portion of the scale. Rather, PAs like these only seem to apply once the object in question has a *sufficient degree* of the relevant degreed property.

Let’s take a closer look at how this works with wetness and dryness. If a towel is completely dry, then it’s not wet to any degree. As Rotstein and Winter have pointed out, even once it’s wet to some (relevant) degree, it’s not automatically wet. Instead, we would describe the towel as ‘moist’ until it’s sufficiently wet to describe it as ‘wet’. For example, imagine having to pull the laundry out of the dryer before the cycle is complete: the laundry can feel not quite dry but also not quite wet. When that happens we’ll describe it as being moist. All of this is represented in the following, more complex **Liminal Model**:




Figure 4: The Liminal Model for Dry/Wet

As this scale seems to have a maximum degree – i.e. it involves a top-closed scale – we note this with a *filled-in black dot* on the right-hand side. In this model, ‘moist’ picks out – what we will call – a

liminal range on the scale: a range *between* being dry and being wet. Towels that fall in this range are wet to some degree, too wet to be considered dry, but simultaneously they are not wet enough to be considered wet (*simpliciter*). It's of note that not all TA/PA scales have a natural language word that refers to the liminal range. For example, there is no single word in English that refers to someone who is *neither naked nor dressed*. But the lack of a single word for this liminal range is no barrier to referring to it; the complex expression 'neither naked nor dressed' does just fine. Liminal Models have not only a liminal range, but a **limited range**, a range of degrees where something may count as *F simpliciter*, while failing to be completely *F*. Again, this is obvious in the case of 'wet': a towel can be wet without being completely wet. Every TA/PA pair that satisfies the Liminal Model has a limited range.²⁴

Note that we have honored the liminal range with about 50% of the visible scale-line. This is an irrelevant geometrical detail. It's perfectly fine if the liminal range in many cases occupies a much smaller proportion of the scale-line. Notice that we used the expression 'completely wet' in connection with the maximum endpoint. It's worth bearing in mind that the expression has two different meanings depending on whether or not the scale has a maximum endpoint. If it does have a maximum endpoint, then 'completely wet' just means 'wet to the maximal degree'. If it does not, then 'completely wet' means 'very wet', as 'completely' can be used as an intensifier.²⁵ If the scale of wetness functions in this second way, then 'completely wet' would mean something more like 'completely imperfect', which does not mean imperfect to the maximum degree possible. There is no such maximum. Rather, 'completely imperfect' just means 'very imperfect'.

Finally, there is one more way that TAs and their antonyms can be associated. Consider, for example, degrees of fullness. On one end of the scale, we have containers that are full *simpliciter*, and on the other end, we have those that are empty *simpliciter* – both of these have conceptual limits. Something is full *simpliciter* iff it is at the point of complete fullness past which a glass cannot become any more full (relative to a contextually specified standard of precision). Something is empty *simpliciter* iff it is at the point of being completely empty past which a glass cannot become any more empty (relative to a contextually specified standard of precision). In this case, both 'empty' and 'full' act as TAs, referring to the endpoint of the underlying scale, like in the following **Bounded Model**:

²⁴A referee raised questions about how 'slightly/somewhat wet' relates to 'moist' and 'wet *simpliciter*'. 'Slightly/somewhat wet' clearly applies to objects that are moist, but, at least in the case of 'somewhat wet', it can also apply to objects that are wet *simpliciter* so long as these objects are close enough to being moist and thus far enough away from being completely wet. In this way, 'somewhat wet' picks out a vague range on the degree of wetness scale that includes the moist range, but can go slightly beyond it. Whether 'slightly wet' can also go beyond the moist range may depend on how small the moist/liminal range is. If the moist/liminal range were to take up only about 15% of the scale, then it strikes us as okay to say of a wet towel that is at around the 16% mark, that it's a slightly wet towel. All of this raises more questions than we can address here.

²⁵ For discussions of the intensifying use of 'completely', see Kennedy (2007: 34-35) and Kennedy and McNally (2005: 354).




Figure 5. Bounded Model for Empty/Full

With the Bounded Model, we have two closed endpoints, one for the full side of the scale and one for the empty side. These two endpoints are separated by a *liminal range*, a region that represents objects that are neither empty nor full. Another potential example of AAs that require the Bounded Model include words like ‘assembled’ and ‘disassembled’. While there are degrees to which a bookshelf can be assembled, it’s only assembled if it’s fully assembled. Likewise, it’s only disassembled if it’s completely so, creating a structure with two endpoints and a liminal range in between.²⁶

5. Modeling Factual Awareness and Its Degrees

Recall the ordering of the distinct epistemic positions of Clueless Carl, Bothered Bart, Seeing Susan, and Knowing Noel in relation to the Bumping fact, i.e. the fact that the noise was caused by a person bumping the wall:




Figure 6. An Awareness Ordering

The primary question of this section centers around how we should model factual awareness *simpliciter* and complete awareness. We will argue for two theses. The first is that the Simple Model is inadequate because factual awareness has a liminal region, pushing us towards thinking of factual

²⁶A reviewer helpfully pointed out that the Bounded Model can still accommodate judgments like ‘somewhat empty’ and ‘somewhat full’ within the liminal range. We can describe a glass as ‘half-empty’ and ‘half-full’, and ‘half-empty’ entails ‘somewhat empty’ while ‘half-full’ entails ‘somewhat full’. So a cup that is filled 50% of the way is both somewhat empty and somewhat full even though it’s not empty or full simpliciter.

awareness either in terms of the Liminal Model or the Bounded Model. The second is that the Liminal Model does a better job of representing the properties associated with factual awareness than the Bounded Model because there is evidence that factual awareness has a limited region, making the Liminal Model the more promising model for factual awareness.

5.1 Is ‘Aware that’ A Relative Gradable Adjective?

Is ‘aware [of the fact] that’ a relative gradable adjective – i.e. an RA? There are several reasons to think it is not. First, were it an RA, then facts about factual awareness would be determined in part by a contextually supplied comparison class. But whether or not Bothered Bart is (un)aware of the Bumping Fact doesn’t depend on some comparison class. Even though *we* have compared Bart to others, as we will argue below, his lack of awareness (*simpliciter*) is not dependent on the comparison we were making. Similarly, as we will argue below, because Seeing Susan is in a good enough situation to *see* the Bumping Fact, she is aware of it (*simpliciter*). No comparisons matter for this determination.

Second, if ‘aware [of the fact] that’ were an RA we would expect it to function like an RA in ‘not at all’ constructions. But it doesn’t. ‘Not at all aware [of the fact] that *p*’ takes us to the zero-degree endpoint of the underlying scale: the point where one has no degree of awareness at all. But consider RAs like ‘tall’, ‘expensive’, and ‘strong’. Expressions like ‘not at all tall’, ‘not at all expensive’ and ‘not at all strong’ don’t take us to the zero-degree endpoint on their underlying scale as they leave it open that something has some degree of height, cost, or strength. Now take their antonym RAs: ‘short’, ‘cheap’, and ‘weak’. ‘Not at all short’ pushes us in the *opposite* direction of the zero-endpoint and towards the contextually salient tall-threshold. The same is true of ‘not at all cheap’ and ‘not at all weak’. This suggests that ‘aware [of the fact] that’ is not an RA.

Third, take *strong* RAs from above, i.e. RAs that don’t gravitate towards the zero-endpoint of their underlying scale. This includes the RAs noted above: ‘expensive’, ‘tall’, and ‘strong’. Strong RAs contrast with the *weak* RAs from above, which do gravitate towards the zero-endpoint: ‘cheap’, ‘short’, and ‘weak’. With strong RAs, expressions of the form ‘somewhat RA’ or ‘RA to some degree’ take us *well beyond* the zero-endpoint on the underlying scale. For example, ‘somewhat tall’ or ‘tall to some degree’ immediately implies that something has a degree of height that is close to, or slightly surpasses, the tall-threshold, which tends to gravitate towards the farther end of the scale. On the height scale, this is never very close to the zero-endpoint. Similarly, ‘somewhat expensive’ or ‘expensive to some degree’ immediately implies that something has a degree of cost that is close to, or slightly surpasses, the contextual expensive-threshold. On the units

of cost scale, this is never very close to the zero-endpoint. So, for strong RAs the expressions ‘somewhat RA’ and ‘RA to some extent’ imply *a significant degree of* the property indicated by the underlying scale. This is further confirmed by the incoherence of statements like: ‘It’s somewhat expensive, but it costs almost nothing’ and ‘He’s somewhat tall, but has very little height.’

Now take *strong* PAs, i.e. PAs that stand on the other side of a liminal range. ‘Somewhat wet [/clothed]’ and ‘wet [/clothed] to some degree’ does not immediately imply that the object in question is wet or clothed to any significant degree. After all, if something is somewhat wet, it may be moist; and if someone is somewhat clothed, they might only have a speedo on. This allows us to coherently say things like ‘the towel is somewhat wet, but it’s almost completely dry’ or that ‘the man is somewhat clothed, but he barely has anything on’. So strong PAs don’t tend to function like strong RAs.

‘Aware [of the fact] that’ functions more like a strong PA than a strong RA. Consider ‘somewhat aware [of the fact] that’ and ‘aware to some extent that’. Neither expression implies a very high degree of awareness, as evidenced by the coherence of the statements like ‘Bart is somewhat [/to some extent] aware of the Bumping Fact, but he’s not aware of it to a significant degree’ and ‘She’s almost completely unaware that he arrived, but she is still somewhat [/to some extent] aware that he did.’ These are not semantically incoherent statements. But we would expect them to be if ‘aware [of the fact] that’ were a strong RA. This is further evidence that ‘aware [of the fact that] that’ is not a relative gradable adjective.

5.2 Possible Models for Factual Awareness

But even if ‘aware [of the fact] that’ is an absolute gradable adjective, that still leaves open the question about whether it should be modeled with the Simple Model, the Bounded Model, or the Liminal Model. To anchor our thinking we present these options visually. First we have the following instance of the Simple Model:




Figure 7. Simple Model For Factual Awareness

On this model, ‘aware [of the fact] that’ is a PA and if factual awareness satisfies this model, there is the following consequence:

C1 If one is aware of a fact that p to any non-zero degree, then one is aware [of the fact] that p .

Now take the Bounded Model:




Figure 8. Bounded Model for Factual Awareness

On this model, ‘aware that’ is a TA equivalent in meaning to ‘completely aware that’, and both expressions refer to the endpoint on the degree of awareness scale. The liminal range between the endpoints, then, captures the thought that between those two endpoints someone may count as not being completely unaware ($\neg CU$) and not being completely aware ($\neg CA$). If factual awareness satisfies this model, there are the following consequences:

¬C1 It is possible to be aware of a fact to some non-zero degree without being aware (*simpliciter*) of it.

C2 Awareness entails complete awareness. That is, it is impossible for one to be aware [of the fact] that p , but fail to be completely aware [of the fact] that p .

Lastly, we have the Liminal Model, which takes ‘aware [of the fact] that’ as a PA that falls on the other side of a liminal range:




Figure 9: Liminal Model for Factual Awareness

The top right brace highlights *the limited range*: a region where one might be aware, but not completely aware. The brace on the bottom left indicates the complement of awareness *simpliciter* – i.e. unawareness *simpliciter* – which includes the liminal range as well as the zero-endpoint.

If factual awareness satisfies the Liminal Model, there are the following consequences:

¬C1 It is possible to be aware of a fact to some degree without being aware (*simpliciter*) of it.

¬C2 Awareness does not entail complete awareness. That is, it is possible for one to be aware [of the fact] that p , but fail to be completely aware of the fact that p .

Both the Bounded Model and the Liminal Model for factual awareness support ¬C1. But unlike the Bounded Model, the Liminal Model entails ¬C2. While these models for factual awareness differ in regard to C1-C3, they agree that a thinker can be unaware to some degree without being completely unaware. Further, given the inter-definability of awareness and ignorance (Silva & Siscoe 2024), we believe this has important implications for understanding both the degreed nature of ignorance as well as the limitations of skeptical arguments. We lack space, however, to discuss either issue.

5.3 Against the Simple Model

Might ‘aware [of the fact] that’ fit the Simple Model? Let’s begin by assessing C1: *is it possible to be aware of the fact that p to some degree, while failing to be aware of the fact that p (simpliciter)?* This question can be difficult to assess apart from concrete cases. When we look at concrete cases a clear answer emerges.

Recall Clueless Carl and Bothered Bart. Clueless Carl is *not at all aware* of the fact that the noise was caused by a person bumping into the wall. But Bothered Bart is *somewhat aware* of that fact.

But should we go on to conclude that Bart is *aware that* the noise was caused by a person bumping into the wall? We think it's hard to take seriously the idea that Bart has awareness (*simpliciter*) of that fact. That is, while Bart is somewhat aware of the Bumping Fact, he remains unaware of it. After all, while Bart knows that a person *might* have caused the noise, he's got great reason to doubt this too since he also knows that this same kind of noise has frequently also been caused by falling supplies. So we find it counterintuitive to regard Bart as being aware of the Bumping Fact, even though we think there is some small degree to which Bart is aware of that fact. Denying this is also odd because, from the first person point of view, Bart himself would likely deny this. That is, Bart, if he is a normal human adult and is relatively cool, calm, and collected, would surely say he's aware of what specific things *might have caused* the bumping sound. But he would not say that he is aware of which one of those things caused the bumping sound.

As further evidence for the claim that Bart is not aware of the Bumping Fact, suppose Bart had a student in his office at the time of the bump. Suppose she asked Bart what caused the sound. If Bart were aware of the fact that it was caused by a person, it would be perfectly appropriate – or at least *not very* inappropriate²⁷ – for Bart to respond by asserting that the sound was caused by a person bumping the wall. But given the details of the case this is a highly inappropriate assertion. After all, Bart has just as much reason to think that the bumping sound was not caused by a person bumping the wall.

Lastly, suppose for *reductio* that Bart were aware that the sound was caused by a person bumping into the wall. If Bart were aware of that fact, then he should be able to become aware of related facts by deduction from that fact. So suppose that Bart reasoned as follows:

P1. The sound was caused by a person bumping the wall.

P2. If the sound was caused by a person bumping the wall, then the sound was not caused by falling supplies that bumped the wall.

C. Therefore, the sound was not caused by falling supplies that bumped the wall.

P2 is true on the supposition that other causes and cases of overdetermination can be ruled out, i.e. cases where both persons and supplies are jointly sufficient causes of the sound. We can just stipulate that this never happens in Bart's case and that Bart knows that this is so. Here's the concern. If we grant that Bart is aware that the sound was caused by a person bumping into the wall, then his deductive abilities will enable him to become aware of the fact *that the sound was not caused by*

²⁷This qualification is owed to the fact that we don't want to take a stance on the norm of assertion. For present purposes it is enough that assertions can be *more or less* appropriate, and asserting facts that one is aware of is much less inappropriate than, say, asserting one's hunches on scant evidence.

falling supplies that bumped the wall. However, this kind of reasoning is an objectionable form of bootstrapping. For recall that Bart has no specific evidence identifying which of the potential causes was the actual cause. So if we afford Bart awareness of P1 just because P1 is true and he's *somewhat* aware of it, then we have to allow for problematic cases of bootstrapping where one becomes aware of other facts that one is clearly not aware of.

In sum, we have strong evidence that an agent can be aware of a fact to such a small degree that they fail to be aware of it (*simpliciter*). This refutes C1. For this reason, verifying \neg C1 is a desideratum for a theory of factual awareness. We have two models that can verify this: the Liminal Model and the Bounded Model.

5.4 Against the Bounded Model

Having eliminated the Simple Model, which of the two remaining models better represents factual awareness? This turns on whether or not the scale of factual awareness has a limited range, and thus allows for a space of cases between awareness (*simpliciter*) and complete awareness. Our conclusion here is modest: the balance of evidence leans towards the existence of a limited range and thus provides at least somewhat more reason to accept \neg C2. This provides non-trivial evidence for the Liminal Model.

The first piece of evidence in favor of \neg C2 involves semantic intuitions. There's just nothing approaching the sound of a contradiction in the following:

(25) ✓ I was aware [of the fact] that my stomach had a serious problem. The cramping was severe. But I was not completely aware [of the fact] that my stomach had a serious problem until the test came back positive for an ulcer.

(26) ✓ She was aware [of the fact] that I was angry with her. She saw my face after that tasteless joke. But, perhaps, she was not completely aware [of the fact] that I was angry with her until after I explained how I felt about her joke.

To our ears, these sound just as good as corresponding claims about wetness, e.g. 'She's wet, but not completely wet. After all, she's only up to her waist in water.' If such claims about awareness can be true without reliance on pragmatic uses of 'aware' and 'completely aware', then \neg C2 follows. Coherent interpretations of these sentences don't seem to require pragmatics.

Next, many have argued that seeing that p entails that one is visually aware that p .²⁸ Many have also argued that one can see that p without knowing that p due to misleading defeaters, i.e. misleading evidence that one does not see that p .²⁹ Bernecker (2010: 78) has argued for the same in the case of remembering that p . If correct, we have the following style of argument in support of $\neg C2$:

- S1. If one sees that p or remembers that p , then one is aware [of the fact] that p .
- S2. It's possible to see that p or remember that p while having misleading defeaters that support the claim that one does not see that p or remember that p .
- S3. Therefore, it's possible for one to be aware [of the fact] that p while having misleading defeaters that support the claim that one is not aware [of the fact] that p . (from S1 and S2)
- S4. If one has misleading defeaters that support the claim that one is not aware [of the fact] that p , then one is not completely aware [of the fact] that p .
- $\neg C2$. Therefore, it's possible for one to be aware [of the fact] that p and yet not be completely aware [of the fact] that p . (from S3 and S4)

Not only does $\neg C2$ follow from S1, S2, and S4, it's intuitive. Seeing Susan is a case where $\neg C2$ is intuitively satisfied: Susan is aware of the Bumping Fact, but she's not completely aware of it because she has significant (misleading) evidence to think that her perceptual capacities have been impaired.

Here's a final argument for $\neg C2$. Again, it will be controversial and requires much more attention than we can devote to it here:

- L1: It's not possible for someone to know that one's lottery ticket is a loser just on the basis of their knowledge that their ticket has a low objective chance of winning in a fair lottery.
- L2: For one to be completely aware [of the fact] that p in a situation c it must be possible for one to know p in c .
- L3: If one's ticket is a loser in a standard lottery case, one is aware of that fact.

L1-L3 imply that lottery-like cases are cases of awareness without complete awareness. L1 is a common position among epistemologists. L2 brings out a conflict between complete awareness and the impossibility of knowledge, and L3 strikes us as intuitive because it's hard for us to think that agents are *ignorant (simpliciter)* in lottery cases. Suppose, for example, I had a ticket in a fair lottery

²⁸ For discussion see Dretske (1993), Huemer (2001), and Silva (2023).

²⁹ McDowell (2002:277), Pritchard (2012: 26), Turri (2010), Schroeder (2021:91), and Silva (2023:51-54).

with long odds and that the winning number had already been drawn. Suppose later I read the news and thus came to know that my ticket is a loser. Upon reading the news, it's not as though I would be surprised by the outcome. The situation is exactly as I expected it to be given my knowledge of the objective chances. So, while I may not have *known* the outcome, it's not obvious that I was *ignorant* of it either. And it seems that if I was not ignorant of the outcome, I was aware of it.³⁰ An interesting consequence of this position is that we can explain the attraction some distinguished epistemologists have felt to regard lottery facts as knowable on the objective chances alone (e.g. Sosa, 2015: 117-124), while not having to go so far as to say that such facts are indeed knowable on the objective chances alone. For factual awareness is entailed by, and comes quite close to, knowledge. But the two remain distinct, yet intimately related factive states. So it's unsurprising that we should find epistemologists defending the idea that we can have knowledge of lottery facts, when those same epistemologists don't carefully distinguish knowledge from factual awareness. In this way, the rejection of the Bounded Model helps provide an explanation of, and a resolution to, a long standing debate in epistemology.

Can the Bounded Model for factual awareness be defended against these arguments? Possibly, for it's often easy to slide from 'somewhat *F*' to '*F simpliciter*'. For example, if you tell me that a towel is somewhat wet, I'm liable to infer that it's wet *simpliciter*. And, arguably, in lottery cases and in the seeing cases (cases where an agent sees that *p*) noted above the target agent is *somewhat aware* of the target fact. But 'somewhat *F*' does not entail '*F simpliciter*'. Rather, when '*F*' is an absolute gradable adjective, 'somewhat *F*' picks out a vague range whose primary limit seems to be that the degree of *F* not be too close to being completely not-*F*. Thus, perhaps, in the lottery and seeing cases above what we are tracking is, in the first, the fact that an agent is *somewhat aware* of a target fact. And the above arguments against the Bounded Model relied not just on the agents being *somewhat aware*, but also that they are *aware simpliciter*. Prospective advocates of the Bounded Model for factual awareness are free to argue that here an illicit transition in thought has occurred: 'somewhat aware' does not entail 'aware *simpliciter*'. And in the lottery and seeing cases above all we

³⁰ For further discussion and defense of factual awareness in lottery cases, see Silva (2023). There the position defended is that mature thinkers are (i) aware of the fact that their ticket is a loser, but (ii) they lack sufficient reason (*ultima facie* propositional justification) to believe it because they know they are not in a position to know it.

can clearly say is that the target agents are somewhat aware, without being completely aware. And this is something that can be true on the Bounded Model.³¹

This defense of the Bounded Model for factual awareness is intriguing. But to our ears, even after the difference between ‘somewhat aware’ and ‘aware *simpliciter*’ has been noted, it remains intuitive to regard agents as being aware *simpliciter* of the target facts in the lottery and seeing cases above. Contrapositively, it remains counterintuitive to us to claim that these agents are not aware *simpliciter* of the target facts when those facts are so strongly supported by the other facts (evidence) they have access to. Indeed, the facts they have access to ensure that it’s highly objectively likely that the target facts are true. This is far from a conclusive argument in defense of the Liminal Model, but it does seem to tip the scales, at least slightly, in favor of it.

Let’s take stock. We have very strong reasons to think that $\neg C1$ is true, and thus to reject the Simple Model for factual awareness. In regard to $\neg C2$, we have presented three distinct considerations that, to our mind, favor the rejection of $C2$ and with it the rejection of the Bounded Model for factual awareness. Admittedly, the case for this is far from conclusive, but it’s noteworthy. Our conclusion, therefore, is modest: the balance of evidence slightly favors the Liminal Model for factual awareness.

6. Concluding Reflections on Knowledge and Awareness

What is the connection between knowledge and awareness? Others have already argued that there exists a minimal connection between knowledge and factual awareness:

Knowledge Requires Awareness. $(K \rightarrow A)$ S knows that p only if S is aware [of the fact] that p .³²

The existing case for $K \rightarrow A$ rests on its direct intuitive appeal and the counterintuitive character of its denial. We think such reasons *also* motivate the idea that:

³¹As a referee pointed out, on an alternative version of the Bounded Model explored by Kennedy (2007:37-38) and Égré (2021:26-28), adjectives that map to a fully bounded scale can be ambiguous between a term that is only used to pick out the maximum point on the scale and a term that can be used to pick out everything except the opposite minimal point on the scale. While this possibility is noteworthy, we already argue in 5.3 that being somewhat aware that p does not entail being aware that p , and we don’t think that there is an interpretation of ‘aware’ where that is the case. Just because Bothered Bart is more aware that the noise was caused by a person bumping into the wall than Clueless Carl, it does not follow that his awareness is strong enough for being aware of that fact *simpliciter*. And we don’t find it compelling to think that there is an alternative disambiguation of ‘aware’ on which Bart is aware of that fact.

³²For defense of this see Dretske (1993), Huemer (2001), Littlejohn (2015), Nagel (2017), and Silva (2023).

Knowledge Requires Complete Awareness. $(K \rightarrow CA)$ S knows that p only if S is completely aware [of the fact] that p .

Like $K \rightarrow A$, $K \rightarrow CA$ strikes us as intuitive. Further, denials of it strike us as counterintuitive. Consider the following:

(27) ?She knows that her father arrived safely, but she's not completely aware [of the fact] that he arrived safely.

Additionally, since not being completely aware is a matter of there being some extent or degree of awareness that is absent, these claims are equivalent to:

(28) #She knows that her father arrived safely, but there's some extent to which she's unaware [of the fact] that he arrived safely.

Both sentences sound very odd to our ears, as bad as denying that a bachelor is unmarried. Nevertheless, perhaps one could get a sensible reading of these by treating 'unaware' as a signal for a lack of *conscious* awareness:

(29) ✓ She knows that her father arrived safely, but there's some extent to which she's not consciously aware [of the fact] that he arrived safely.

This sentence is perfectly fine. But it's fine because it invites a contrast between merely dispositional knowledge (a relation to facts one can have even while asleep)³³ and conscious awareness of facts (a relation to facts one cannot have while asleep). But not all awareness is conscious. As we have discussed above: if you've learned and not forgotten that $2+2=4$, then you're aware that $2+2=4$. And this is true of you even if you're asleep or not at all consciously thinking about that equation. So the relevant claim to consider is something like the following:

(30) #Even though she's asleep and unconscious, she knows that her father arrived safely. But there's some extent to which she's not aware of the fact that he arrived safely.

³³Compare Audi (1994).

And this is problematic. The reason, we propose, is this: knowing that p requires complete awareness. Irrespective of whether or not we opt for the Bounded Model or the Liminal Model, it's a consequence of $K \rightarrow CA$ that one can have some degree of awareness of a fact without knowing it, just as one can have some degree of awareness of a fact without being completely aware of it.

While the case of $(K \rightarrow CA)$ is reasonably strong or at least noteworthy, what should we think of its converse? That is, what should we make of:

Complete Awareness Entails Knowledge. $(CA \rightarrow K)$ If S is completely aware of the fact that p , then S knows that p .

Should $CA \rightarrow K$ turn out true, we'd have the makings for a novel *structural* theory of knowledge: knowledge *just is* complete awareness. This would be a structural theory of knowledge in the sense that it provides a reductive definition of knowledge in terms of a maximum degree of factual awareness, without going on to say anything about what factual awareness is. In this way, it's a theory of knowledge that is consistent with many standing debates about its relation to belief, justification, reasons, safety, sensitivity, reliability, proper function, and so forth.

We are attracted to this. But we also have reservations. For example, many think that knowledge requires belief, while factual awareness (*simpliciter*) does not require belief.³⁴ For this reason, we think it's somewhat more promising to float the idea that *being in a position to know* just is *being completely aware of a fact*. But this turns on:

Complete Awareness Entails Position to Know. $(CA \rightarrow PK)$ If S is completely aware of the fact that p , then S is in a position to know that p .

Being in a position to know is, on all counts, not belief entailing. So no problems in connection with the belief requirement on knowledge arise for $CA \rightarrow K$.

However, other problems may arise. For example, McGlynn (2014: 63) and Silva (2023: 54-58) have argued that in fake barn cases (and other cases of environmental luck) one can be aware of the fact *that one is looking at a barn* even though one is not in a position to know this. If that's right,

³⁴See Silva (2023) and Silva and Siscoe (2024). Also of relevance is the work of those who think we can see that p or remember that p without believing that p (McDowell, 2002; Turri, 2010; Bernecker, 2010; Pritchard, 2012; Schroeder, 2021; Silva, 2023). For, arguably, if one sees that p or remembers that p , then one is aware of the fact that p .

then we have to ask whether or not complete awareness is compromised by luck in the same way that knowledge is. If not, then there is some hope for holding that being in a position to know just is complete awareness. But if it is, then such a hope will likely erode. As you may have begun to notice, the issues here are complex, subtle, and deserve more discussion than we can here provide.

In the end, our hope is that we've provided the foundations necessary for readers to effectively and efficiently bring factual awareness and its degrees to bear on issues in epistemology in ways that have so far been neglected.³⁵

³⁵ We are grateful to a referee for pointing out that '*S* vaguely knows that *p*' seems to entail '*S* is somewhat aware [of the fact] that *p*'. We think there are a wide and interesting range of questions to explore about vague knowing and its relation to: being somewhat aware, being aware *simpliciter*, being in a position to know, and knowing. However, we lack space to explore these questions here.

References

- Audi, R. (1994). Dispositional beliefs and dispositions to believe. *Noûs* 28.4, 419-34. doi: 10.2307/2215473
- Baumann, P. (2016). *Epistemic Contextualism: A Defense*. Oxford University Press.
- Beddor, B. (2020). New work for certainty. *Philosophers' Imprint* 20.8, 1-25,
- Bengson, J., & Moffett, M. (2011). Non-propositional intellectualism. In J. Bengson & M. Moffett (Eds.), *Knowing How: Essays on Knowledge, Mind, and Action* (pp. 161-195). Oxford University Press.
- Bernecker, S. (2010). *Memory: A Philosophical Study*. Oxford University Press.
- Block, N. (2016). The anna parenina principle and skepticism about unconscious perception. *Philosophy and Phenomenological Research*, 93.2, 452–459. doi: 10.1111/phpr.12258
- Blome-Tillmann, M. (2017). Gradability and knowledge. In J. Ichikawa (Ed.), *The Routledge Handbook of Epistemic Contextualism* (pp. 348-357). London: Routledge.
- Blome-Tillmann, M. (2022). *The Semantics of Knowledge Attributions*. Oxford University Press.
- Brogaard, B. (2016). Ignorance and incompetence: Linguistic considerations. In R. Peels & M. Blaauw (Eds.), *The Epistemic Dimensions of Ignorance* (pp. 57-80). Cambridge University Press.
- Burnett, H. (2014). A delineation solution to the puzzles of absolute adjectives. *Linguistics and Philosophy*, 37, 1-39. doi: 10.1007/s10988-014-9145-9
- Burnett, H. (2017). *Gradability in Natural Language: Logical and Grammatical Foundations*. Oxford University Press.
- de Canson, C. (2024). The nature of awareness growth. *Philosophical Review*, 133.1, 1-32. doi: 10.1215/00318108-10880419
- Chalmers, D. (1996). *The Conscious Mind*. Oxford University Press.
- Crane, T. (2012). Review: Tye on acquaintance and the problem of consciousness. *Philosophy and Phenomenological Research*, 84.1, 190–198. doi: 10.1111/j.1933-1592.2011.00557.x
- DeRose, K. (2011). *The Case for Contextualism: Knowledge, Skepticism, and Context*. Oxford University Press.
- Dretske, F. (1981). The pragmatic dimension of knowledge. *Philosophical Studies*, 40.3, 363-379. doi: 10.1007/BF00646423
- Dretske, F. (1993). Conscious experience. *Mind*, 102.406, 263-283. doi: 10.1093/mind/102.406.263
- Dutant, J. (2007). The limited gradability of 'knows'. Accessed April 8, 2023, from <http://julien.dutant.free.fr/papers/JDutant_TheLimitedGradabilityOfKnow.pdf>
- Egré, P. (2021). Half-truths and the .iar. In C. Nicolai and J. Stern (Eds.), *Modes of Truth: The Unified Approach to Truth, Modality, and Paradox* (pp. 18-40). Routledge.
- Fassio, D. & Logins, A. (2023). Justification and gradability. *Philosophical Studies* 180.7, 2051-2077. doi: 10.1007/s11098-023-01945-3
- Fazekas, P. & Morten O. (2018). A multi-factor account of degrees of awareness. *Cognitive Science*, 42, 1833–1859. doi: 10.1111/cogs.12478
- Frege, G. (1977). Thoughts. In P. Geach (Eds.), P. Geach and R. Stoothoff (Trans.), *Logical Investigations*. Oxford: Basil Blackwell.
- Hawthorne, J., & Logins, A. (2021). Graded epistemic justification. *Philosophical Studies*, 178.6, 1845-1858. doi: 10.1007/s11098-020-01512-0
- Hetherington, S. (2001). *Good knowledge, Bad knowledge: On Two Dogmas of Epistemology*. Oxford University Press.
- Hetherington, S. (2011). *How to Know: A Practicalist Conception of Knowledge*. Wiley-Blackwell.
- Hills, A. (2016). Understanding why. *Noûs*, 50.4, 661-688. doi: 10.1111/nous.12092

- Huemer, M. (2001). *Skepticism and the Veil of Perception*. Lanham: Rowman & Littlefield.
- Kennedy, C. (2007). Vagueness and grammar: The semantics of relative and absolute gradable adjectives. *Linguistics and Philosophy*, 30.1, 1-45. doi: 10.1007/s10988-006-9008-0
- Kennedy, C. & McNally, L. (2005). Scale structure, degree modification, and the semantics of gradable predicates. *Language*, 81.2, 345-381. doi: 10.1353/lan.2005.0071
- Kennedy, C. & McNally, L. (2012). Degree vs. manner *well*: A case study in selective binding. In J. Pustejovsky, P. Bouillon, H. Isahara, K. Kanzaki, and C. Lee (Eds.), *Advances in Generative Lexicon Theory* (pp. 247-262). Dordrecht: Springer Netherlands.
- Knowles, J. (1974). The degree adverbial. *Journal of English Linguistics*, 8, 21–31. doi: 10.1177/007542427400800103
- Krifka, M. (2002). Be brief and vague! And how bidirectional optimality theory allows for verbosity and precision. In D. Restle & D. Zaefferer (Eds.), *Sounds and Systems, Studies in Structure and Change: A Festschrift for Theo Vennemann* (pp. 439–458). Berlin: Mouton de Gruyter.
- Krifka, M. (2007). Approximate interpretations of number words: A case of strategic communication. In G. Bouma, I. Kramer, and J. Zwarts (Eds.), *Cognitive Foundations of Interpretation*. Amsterdam: Koninklijke Nederlandse Akademie van Wetenschappen.
- Kvanvig, J. (2003). *The Value of Knowledge and the Pursuit of Understanding*. Cambridge University Press.
- Lakoff, G. (1973). Hedges: A study in meaning criteria and the logic of fuzzy concepts. *Journal of Philosophical Logic*, 2, 458-508. doi: 10.1007/BF00262952
- Lai, C. (2021). Against epistemic absolutism. *Synthese*, 199.1-2, 3945-3967. doi: 10.1007/s11229-020-02964-x
- Lassiter, D. (2017). *Graded Modality*. Oxford University Press.
- Lee, A. (2022). Degrees of consciousness. *Nous*, 57.3, 553-575. doi: 10.1111/nous.12421
- Lewis, D. (1979). Scorekeeping in a language game. *Journal of Philosophical Logic*, 8.1, 339–359. doi: 10.1007/BF00258436
- Littlejohn, C. (2015). Knowledge and awareness. *Analysis*, 75.4, 596–603. doi: 10.1093/analys/anv051
- Littlejohn, C. (2017). How and why knowledge is first. In J.A. Carter, E.C. Gordon, and B. Jarvis (Eds.), *Knowledge First* (pp. 19-45). Oxford University Press.
- McDowell, J. (2002). Responses. In N.H. Smith (Ed.), *Reading McDowell on Mind and World* (pp. 269–305). Routledge.
- McGlynn, A. (2014). *Knowledge First?* New York, NY: Palgrave Macmillian.
- Moore, G.E. (1959). Certainty. In *Philosophical Papers*. London: Allen and Unwin.
- Nagel, J. (2013). Knowledge as a mental state. *Oxford Studies in Epistemology*, 4, 272-308. doi: 10.1093/acprof:oso/9780199672707.003.0010
- Nagel, J. (2017). Factive and nonfactive mental state attribution. *Mind and Language*, 32.5, 525–544. doi: 10.1111/mila.12157
- Pappas, G. (2017). Review of Stephen Hetherington’s “How To Know: A Practicalist Conception of Knowledge.” *Dialectica*, 71.2, 309-311. doi: 10.1111/1746-8361.12181
- Pavese, C. (2017). Know-how and gradability. *The Philosophical Review* 126.3:345-383. doi: 10.1215/00318108-3878493
- Pinkal, M. (1995). *Logic and Lexicon*. Dordrecht: Kluwer.
- Prinz, J.J. (2010). When is perception conscious? In B. Nanay (Ed.), *Perceiving the World Oxford* (pp. 310-332). Oxford.
- Pritchard, D. (2012). *Epistemological Disjunctivism*. Oxford University Press.
- Roberts, C. (2009). Know-how: A compositional approach. In E. Hinrichs & J. Nerbonne (Eds.), *Theory and Evidence in Semantics* (pp. 183-213). CSLI Press.

- Rotstein, C. & Winter, Y. (2004). Total adjectives vs. partial adjectives: Scale structure and higher-order modifiers. *Natural Language Semantics*, 12, 259-288. doi: 10.1023/B:NALS.0000034517.56898.9a
- Rusiecki, J. (1985). *On Adjectives and Comparisons in English*. New York: Longman Linguistics Library.
- Ryle, G. (1949). *The Concept of Mind*. University of Chicago Press.
- Sadock, J. (1977). Truth and approximations. In K. Whistler Jr., R.D. Van Valin, C. Chiarello, J.J. Jaeger, & M. Petrucci (Eds.), *Proceedings of the Third Annual Meeting of the Berkeley Linguistics Society*. Berkeley Linguistics Society.
- Sauerland, U. & Stateva, P. (2007). Scalar vs. epistemic vagueness: Evidence from approximators. In T. Friedman & M. Gibson (Eds.), *Proceedings of the Seventeenth Semantics and Linguistic Theory Conference*. Ed. Ithaca, NY: CLC Publishing.
- Schaffer, J. & Szabó, Z.. (2014). Epistemic comparativism: A contextualist semantics for knowledge ascriptions. *Philosophical Studies*, 168.2, 491-543. doi: 10.1007/s11098-013-0141-7
- Schroeder, M. (2021). *Reasons First*. Oxford University Press.
- Silva, P. & Bernecker, S. (2023). Evidence, reasons, and knowledge in the reasons-first program. *Philosophical Studies*, 181.2, 617-625. doi: 10.1007/s11098-023-02044-z
- Silva Jr., P. (2023). *Awareness and the Substructure of Knowledge*. Oxford University Press.
- Silva Jr., P. (Forthcoming). Awareness. In K. Sylvan (Ed.), *Blackwell Companion to Epistemology*.
- Silva Jr., P. & Siscoe, R.W. (2024). Ignorance and awareness. *Noûs*, 58.1, 225-243. doi: 10.1111/nous.12450
- Siscoe, R.W. (2021a). Belief, rational and justified. *Mind*, 130.517, 59-83. doi: 10.1093/mind/fzaa021
- Siscoe, R.W. (2021b). Does being rational require being ideally rational? *Philosophical Topics*, 49.2, 245-266. doi: 10.5840/philtopics202149224
- Siscoe, R.W. (2022a). Real and ideal rationality. *Philosophical Studies*, 179, 879–910. doi: 10.1007/s11098-021-01698-x
- Siscoe, R.W. (2022b). Rational supererogation and epistemic permissivism. *Philosophical Studies*, 179, 571–591. doi: 10.1007/s11098-021-01671-8
- Siscoe, R.W. (2023). Being rational enough: Maximizing, satisficing, and degrees of rationality. *Australasian Journal of Philosophy*, 101.1, 111-127. doi: 10.1080/00048402.2021.1942940
- Siscoe, R.W. (2024). Knowledge, true belief, and the gradability of ignorance. *Philosophical Studies*, 181, 893-916. doi: 10.1007/s11098-024-02119-5
- Sorensen, R. (1991). Rationality as an absolute concept. *Philosophy*, 66.258, 473-486. doi: 10.1017/s0031819100065128
- Sosa, E. (2001). Human knowledge, animal and reflective. *Philosophical Studies*, 106, 193-196. doi: 10.1023/A:1013300923093
- Sosa, E. (2011). *Knowing Full Well*. Princeton University Press.
- Sosa, E. (2015). *Judgment and Agency*. Oxford University Press.
- Stanley, J. 2004. On the linguistic basis for contextualism. *Philosophical Studies*, 119.1-2, 119-146. doi: 10.1023/B:PHIL.0000029353.14987.34
- Stanley, J. (2005). *Knowledge and Practical Interests*. Oxford University Press.
- Stanley, J. (2008). Knowledge and certainty. *Philosophical Issues*, 18.1, 35-57. doi: 10.1111/j.1533-6077.2008.00136.x
- Stanley, J. (2011). *Know How*. Oxford University Press.
- Steele, K. & Stefánsson, H.O. 2021. Belief revision for growing awareness. *Mind*, 130.520, 1207–1232. doi: 10.1093/mind/fzaa056
- Strawson, P. F. (1998). Reply to John Searle. In L.E. Hahn (Ed.) *The Philosophy of P. F. Strawson*. La Salle, IL: Open Court.

- Sylvan, K. (2018). Knowledge as a non-normative relation. *Philosophy and Phenomenological Research*, 97.1, 190-222. doi: 10.1111/phpr.12371
- Turri, J. (2010). Does perceiving entail knowing? *Theoria*, 76.3, 197–206. doi: 10.1111/j.1755-2567.2010.01065.x
- Turri, J. (2015). Unreliable knowledge. *Philosophy and Phenomenological Research* 90.3: 529–545. doi: 10.1111/phpr.12064
- Unger, P. (1975). *Ignorance: A Case for Skepticism*. Oxford University Press.