



Knowing Falsely: the Non-factive Project

Adam Michael Bricker¹

Received: 21 October 2020 / Accepted: 10 May 2021 / Published online: 3 July 2021
© The Author(s) 2021

Abstract

Quite likely the most sacrosanct principle in epistemology, it is near-universally accepted that knowledge is factive: knowing that p entails p . Recently, however, Bricker, Buckwalter, and Turri have all argued that we can and often do know approximations that are strictly speaking false. My goal with this paper is to advance this nascent non-factive project in two key ways. First, I provide a critical review of these recent arguments against the factivity of knowledge, allowing us to observe that elements of these arguments mutually reinforce respective weaknesses, thereby offering the non-factive project a much stronger foundation than when these arguments were isolated. Next, I argue tentatively in favor of Bricker's truthlikeness framework over the representational adequacy account favored by Buckwalter and Turri. Taken together, while none of this constitutes a knock-down argument against factivity, it does allow us to quiet some of the more immediate worries surrounding the non-factive project.

Keywords Epistemology · Knowledge · Analysis · Factivity · Truth condition

1 Introduction: the Factivity of Knowledge

At the heart of this paper is a remarkably simple question: Is knowledge factive? That is, does S knowing that p require the truth of p ? The almost-universal attitude in epistemology is that this simple question invites an equally simple answer: Yes, knowledge is in fact factive. Indeed, it is difficult to identify a less controversial principle in epistemology than the factivity of knowledge. With origins dating to Plato's *Meno* and *Theaetetus*, the textbook view is that knowledge is simply, obviously factive (see e.g. Audi, 2010, 246; BonJour, 2010, 32; Pritchard, 2017, 4), an idea we see widely repeated in contemporary epistemology (see e.g. Armstrong, 1973, 137; Davidson, 1988, 177; Pritchard & Kelp, 2010, 330; Hannon, 2013). Indeed, this

✉ Adam Michael Bricker
adam.michael.bricker@gmail.com

¹ Cologne Center for Contemporary Epistemology and the Kantian Tradition, Department of Philosophy, University of Cologne, Albertus-Magnus-Platz, 50923 Köln, Germany

idea is so widespread that we even see it reflected in the work of the epistemologists whose non-factive accounts are the subject of the present paper (Turri, 2011; Bricker, 2019, §2; Buckwalter, 2014). The point here is simply that epistemologists have long agreed not only that knowledge is factive, but that whether knowledge is factive doesn't even constitute an epistemologically significant question.

Until very recently, the unanimous epistemological agreement around factivity carried with it only a very small number of scattered asterisks, and I've struggled to identify more than four. First, likely the most familiar asterisk comes from Hazlett, who has argued that the ordinary concept governing use of the verb "knows" is non-factive (2010). However, not only did Hazlett explicitly reject the idea that this gave us reason to think that the epistemologically interesting concept of knowledge was non-factive, but subsequent experimental evidence has undermined his account of how "knows" is used in the first place (Buckwalter, 2014). Second, Niiniluoto favorably considered an explicitly non-factive formulation of knowledge on the basis that the knowledge produced by science is often "not true in a strict sense" (1999, 84). However, as Niiniluoto's primary concern was not the analysis of knowledge, but instead scientific realism, he devotes very little time to arguing for abandoning factivity, and it isn't entirely clear the extent to which he endorses the non-factive formulation he considers. Perhaps for this reason, despite appearing in a monograph with considerable impact in the philosophy of science,¹ Niiniluoto's account has attracted remarkably little attention from epistemologists working on the analysis of knowledge. Next, the *Stanford Encyclopedia* article on the analysis of knowledge notes a quite similar sort of objection regarding knowledge of false scientific theories, which is attributed to Nicholas Maxwell (Ichikawa & Steup, 2018, footnote 3). However, as such an objection doesn't seem to appear anywhere in print, I can't really consider it here. Finally, perhaps the only bona fide non-factive account of knowledge within the contemporary analytic project came from Ackermann's somewhat obscure 1972 monograph *Belief and Knowledge*, which maintains that our knowledge in certain cases of "everyday fact" is inconsistent with the truth condition (115). However, in what little attention this argument has attracted, it has been roundly criticized, both contemporarily (Annis, 1974; Shope, 1983, 197–199) and more recently (Bricker, 2018, 8–11). In short, whatever asterisks there may have been on the unanimity of the epistemological assent to factivity, we might understand such asterisks to be closer to negligible than substantive.

Given this longstanding consensus that knowledge is factive, the recent emergence of an epistemological program questioning factivity represents a particularly interesting development in philosophy. First, Bricker's 2018 doctoral monograph presented a comprehensive argument against the factivity of knowledge, drawing heavily from empirical features of human perceptual systems. Shortly thereafter, Buckwalter and Turri published a pair of papers arguing the same on both theoretical and experimental grounds (2020a, b). More interesting still is that despite by all indications having developed independently from one another, these two accounts share a number of major structural similarities. Most notably, both maintain that (i)

¹ As of 6 April 2021, a Google Scholar search returns 911 citations of *Critical Scientific Realism*.

the ubiquity of approximately-true-yet-false beliefs in daily life creates a skeptical problem for factivity and (ii) intuitive judgements about knowledge sometimes don't align with factivity in such cases of merely approximate truth. While too early to tell if non-factivity is in the air, or instead if this is just a coincidence, all this raises the question of just how viable this emerging non-factive project is. While here I will by no means present anything like a definitive argument against the factivity of knowledge, my aim is nonetheless to advocate for the viability of the non-factive project, bolstering it in a few key respects.

Our present task of bolstering the viability of the non-factive project might be divided into two key elements. First, there is the task of addressing a couple of extant shortcomings present in the arguments of both Bricker as well as Buckwalter and Turri. This occupies the first two sections of this paper, in which I survey the two major arguments against factivity present in both the accounts of Bricker and Buckwalter/Turri—the argument from skepticism (§2) and the argument from intuitive judgements about knowledge (§3). In both cases, we might observe that elements of one argument reinforce weaknesses in the other, thereby resulting in stronger challenges to factivity than were originally present in either. Most importantly, for the skeptical argument, Bricker's empirical argument from specific cases of cognitive approximation directly addresses one of the major objections left open by the Buckwalter and Turri argument from naïve approximation—i.e., that it's unclear whether belief content in cases of naïve approximation is so finely grained that such approximately true beliefs are indeed false (§2). Conversely, for the argument from intuitive judgements about knowledge, experimental results from Buckwalter and Turri confirm Bricker's entirely armchair speculation that we sometimes attribute knowledge in cases of merely approximate truth, additionally calling into question Bricker's strict commitment to a certain species of inaccessibility in such cases (§3). Next, after observing how this interaction between arguments mutually reinforces some of the more glaring of their respective weaknesses, I then seek to advance the non-factive project further by arguing in favor of the truthlikeness framework advanced by Bricker over the representational adequacy framework of Buckwalter and Turri (§4). Put simply, the truthlikeness approach, on which proximity to truth takes theoretical priority over pragmatic matters, is much more successful at preserving the fundamental connection between knowledge and truth. Accordingly, it offers the more promising starting point for the future development of the non-factive project. Finally, I close by reflecting on where this leaves the project, along with the challenges that lie ahead (§5).

2 The Arguments from Skepticism

The following two sections present a critical review of the major arguments against the factivity of knowledge identified by both Bricker and Buckwalter/Turri. The focus of the present section will be their arguments from skepticism, which are prominently featured in both accounts and remarkably similar in overall structure. One crucial point of separation, however, centers around an objection left open by Buckwalter and Turri's argument from cases of naïve approximation—roughly, that

such beliefs are too coarsely grained to be false in the way imagined by the non-factive project. As Bricker's argument from cases of cognitive approximation was largely built around eliminating similar concerns, it provides us with the exact tools required to overcome objections related to the grain of approximative belief content. Ultimately, doing so provides us with a deeper understanding of the nature of the skeptical challenge posed by factivity. This discussion of the arguments from skepticism will begin with the outline of the argument from naïve approximation presented by Buckwalter and Turri, along with the objection regarding the grain of belief content. After this, I'll then discuss how the argument from cognitive approximation developed by Bricker differs from the argument from naïve approximation in a way that allows us to meet this objection.

Let's begin with what I'm referring to here as the argument from naïve approximation. Perhaps the most immediately striking feature of this argument is that despite opening up a novel line of resistance to a foundational epistemological principle, the skeptical challenge it outlines is remarkably straightforward. Buckwalter and Turri set up the argument by noting that we rely on a wide range of approximations in daily life, citing such examples as (1) "a mile equals 1.6 km," (2) "my mailbox is 2 ft from the street," (3) "16 oz of water was added to the recipe," (4) "two blue shirts are the same shade," (5) "the current temperature is 75° F outside," and (6) "the current time is 9:03" (2020a, 94). As these examples illustrate, the cases utilized by this argument are characterized by an intuitive understanding of ordinary, approximative language. The approximative nature of such representations is taken to be largely self-evident, attracting little argumentation, and here I will not challenge this treatment. After all, when, for example, we say things like, "The time is 9:03," this is merely an approximation. It is almost certainly not precisely 9:03, but instead something close to it.

Buckwalter and Turri further maintain that "these representations closely approximate the truth, but they are not strictly true" (2020a, 94). While I will return to this question later in the section, at least initially we might note that this seems plausible. After all, if it isn't actually 75 °F outside, but instead 75.1 °F, a representation on which it is 75 °F outside won't strictly match reality. Accordingly, Buckwalter and Turri contend that such cases, when coupled with factivity, pose a skeptical problem. These approximate representations can never constitute the content of known beliefs, as they will never be true. However, given the ubiquity of such approximations, this then means that "much of what we ordinarily take ourselves to know" will fail to actually count as knowledge (2020a, 95). Herein lies the crux of the skeptical challenge identified in the argument from naïve approximation. Since many of our everyday beliefs, which we ordinarily take to constitute knowledge, are in fact only approximately true, such beliefs cannot be known, given that knowledge is factive. Again, the structure of the argument is remarkably straightforward, and there is little more to say about the argument itself at this point.

Instead, I want to discuss one of the potential objections flagged by Buckwalter and Turri (2020a, 95–96)—the skeptical argument they outline rests on the relevant representational content being quite fine grained. If the content of our beliefs in cases of naïve approximation turned out to be imprecise to the point that we will not normally have almost-true-but-false beliefs, this would undermine the proposal

that these approximations are actually false. For example, when we believe it is 75 °F outside, if this belief content is imprecise enough to also include 75.1 °F, the belief of course won't be false when it is in fact 75.1 °F outside. As this seems like a natural way to understand the content of these sorts of beliefs, we might understand that it constitutes a serious objection to the argument from naïve approximation. In the end, Buckwalter and Turri don't offer a direct refutation of this objection. Instead, they maintain that it is ultimately an empirical matter, one which has yet to be addressed by empirical research one way or the other. Accordingly, while they seek to diffuse to the objection, they ultimately leave it as something like an open problem. In what follows, I aim to show that there is actually a good deal of empirical research directly relevant to this question of content grain, thereby allowing us to close this problem and strengthen the case against factivity. However, before such time as the grain problem might be dismissed, it is clear that this constitutes a significant limitation for the argument from naïve approximation.

Even before specifying this open problem, however, one might get the sense that more work needs to be done to establish the plausibility of the putative skeptical challenge posed by factivity. As sympathetic as I am to the sort of argument that Buckwalter and Turri present, I must admit that even I don't find it especially convincing in its present form, and I suspect that I am by no means an outlier in this regard. Accordingly, the rest of this section will be devoted towards discussing how cases of cognitive approximation can avoid the grain problem identified by Buckwalter and Turri, providing us with the means to describe a more potent version of the skeptical challenge. Ultimately, the key point here is that the anti-skeptical motivation behind the non-factive project runs deeper than may have been suggested by the argument from naïve approximation alone.

Let's now turn to what here I'm referring to as the argument from cognitive approximation.² Much like we saw in the case of naïve approximation, the argument from cognitive approximation maintains that many of our everyday beliefs, which we ordinarily consider to constitute knowledge, are only approximately true and strictly speaking false. Therefore, if knowledge is factive, then these everyday beliefs cannot actually constitute knowledge, thereby creating a skeptical problem for factivity. Moreover, also like in the case of naïve approximation, the argument from cognitive approximation concludes on this basis that knowledge isn't strictly factive and that approximately-true-yet-false beliefs can still be known (see e.g. Bricker, 2018, ch. 1 §1.2). In short, at least with respect to the broad strokes of their respective arguments, Bricker and Buckwalter/Turri are in agreement about dropping factivity for anti-skeptical reasons.

However, despite these broad structural similarities, this argument eschews much of what characterizes the naïve approximations discussed above. First, we might observe that not only does Bricker not rely on the sort of naïve understanding of

² Note that while all the elements of this argument have been previously advanced by Bricker (2018), their presentation is much more complicated than is conveyed in my review, and a number of extraneous components are omitted entirely. Here, my aim has been to present the integral components of the argument in a stronger and more straightforward configuration than they were previously.

approximative language at the heart of the argument from naïve approximation, but he explicitly rejects the idea that the beliefs associated with such language are even false in the first place:

After all, we seem to have no trouble attributing knowledge for S's belief that a cup is at the edge of the table, even if it is in fact offset a single micron from the edge, and likely for much larger offsets as well. Doesn't this suggest that factivity is weaker than the above formulation would suggest? While such sentiment is understandable, I don't think this is correct. The approximation here is a matter of language, not belief content. Were we to probe our agent's belief that the cup is on the edge of the table, it's unlikely we'd find she actually believed the cup was on the exact edge of the table. The English sentence, "The cup is on the edge of the table," is entirely compatible with the belief that the cup is in the region of the edge, but not the exact edge, and I expect this is normally the belief we hold when uttering such a sentence. Moreover, if in the course of probing the belief's contents it was instead revealed that S actually did believe that the cup was on the exact edge of the table, with not a single micron of offset or overhang, we would then rescind our judgement that she knew that the cup was on the edge of the table, and explain this judgement by appealing to the factivity condition. (2018, 50)

I would add further that this same idea also explains why we don't consider something like, "Actually, it's 75.1°" to be a legitimate correction to S's asserting that it is 75 °F outside. Instead, we would dismiss it as somewhere between a joke and a nuisance. S of course doesn't believe that it is exactly 75° F outside, but just that it is approximately 75° F outside. In short, here, it is far more likely that at least many of the beliefs that form the basis of the argument from naïve approximation are not approximately true, but instead truly approximate. Not only does this sentiment sharpen the grain problem discussed above—after all, the point here is just that these beliefs don't seem as precise as Buckwalter and Turri assume—but it also might explain some of the argument from naïve approximation's seeming dearth of convincingness. The non-factive project, it would appear, needs a stronger class of approximations to motivate the argument from skepticism.

As it happens, cognitive approximations (which we might roughly characterize as representational content closely associated with specific cognitive processes or systems, the approximative nature of which is identified and described empirically) let us do just that. In order to understand why at least some cases of cognitive approximation fit the bill, let's begin with the "recipe" Bricker offers for the identification of cases of cognitive approximation that can support an argument from skepticism against factivity (2018, ch.5 §1): First, one needs to identify "appropriate mental content," which is characterized here as "mental content that might misrepresent the world, but only slightly" (2018, 98). While something resembling this step is also implicit in the way that the argument from naïve approximation begins with (putative) approximately true representations, the present method notably differs by explicitly seeking out empirical evidence to facilitate this identification of slight misrepresentation. As opposed to the armchair approach taken by Buckwalter and Turri, Bricker relies heavily on empirical evidence to establish that certain mental

content is in fact only approximately true. As we'll see shortly, this is precisely what allows such cases of cognitive approximation to meet the grain objection encountered by the argument from naïve approximation. After this initial empirical task of identifying mental content, the process employed by Bricker then becomes increasingly conceptual. One needs to identify in turn: (i) certain behaviors associated with that mental content; (ii) beliefs we might attribute on the basis of that behavior; and finally (iii) any such beliefs we might consider knowledge for, e.g., anti-skeptical reasons. The idea is that if there are any such cases that pose a skeptical problem for factivity, this represents the bottom-up, empirically grounded means of locating them.

While Bricker applies this method to a number of different empirical cases across a range of cognitive processes and systems, here I will only focus on one, which I take to be the most convincing and straightforward—foveal bias, on which the representation of visual space is systematically biased towards the center of vision (the fovea).³ An important early demonstration of foveal bias comes from a 2001 study authored by Sheth and Shimojo, in which participants were asked to use a computer mouse to point to the on-screen location of briefly presented stimuli. Despite the task itself being quite simple, it revealed not only that participants struggled to accurately represent the stimuli's precise locations in visual memory, but that they "systematically mislocalized the target closer to the center of gaze" (329). That is, their representations of target location, on which they based their subsequent actions, misrepresented targets as being closer to their center of vision than they actually were. To be clear, mislocalization wasn't especially severe and typically ranged in the order of magnitude of single degrees (331), but it was nonetheless easily observable in a laboratory setting. A large number of subsequent studies have observed this sort of localization bias not just towards the center of the visual field (Bocianski et al., 2008; Brenner et al., 2008; Fortenbaugh & Robertson, 2011; Kerzel, 2002; Odegaard et al., 2015; Uddin et al., 2005; Zimmermann et al., 2013), but a similar effect has also been observed in relation to foci of attention (e.g. Ono & Watanabe, 2011; Suzuki & Cavanagh, 1997). Crucially, however, foveal bias does not reflect some illusionary feature of anomalous stimuli, but instead a general limitation of our visual systems. As put by Brenner et al., "Its origin lies in the way in which visual information is processed within the eye and brain" (2008, 2). While there is disagreement regarding the precise mechanisms governing the effect (compare Brenner et al., 2008; Odegaard et al., 2015; Sheth & Shimojo, 2001), we have little reason to doubt that these mechanisms frequently result in non-zero foveal bias during daily life. In short, the available empirical evidence strongly suggests that our representations of objects' locations are systematically biased towards our center of vision, however slightly, and therefore not strictly true.

With this empirical background in place, following Bricker's "recipe," we are then in a position to ask which behaviors are associated with our representations of

³ Note that while this largely recapitulates Bricker's discussion of foveal bias (2018, 101–103), I've also included a number of additional citations, framed discussion of the underlying mechanism differently, and generally opted for a different presentational structure.

objects in visual space, which might then be slightly biased towards the center of vision.⁴ In the specific case of foveal bias, this task is simplified both by the behavioral evidence gained in studying the effect and the general, systemic nature of the misrepresentation. Put simply, there are a large class of both motor (e.g. pointing, reaching) and intellectual (e.g. reporting on an object's location, using it in reasoning) behaviors associated with our representations of objects in visual space, and these representations will generally be susceptible to foveal bias.⁵ Crucially, this means that the representations guiding these behaviors will often, however slightly, misrepresent the location of objects in visual space.

Next, we might now identify at least one major class of beliefs that we might attribute on the basis of behaviors associated with these representations—beliefs about an object's location.⁶ Clearly, when, on the basis of her perceptions or visual memory, S points to/reaches towards/reports on the location of/etc. some object, we have no problem concluding that she has some belief about the location of that object, a belief which we likely would express with a demonstrative, "S believes that the object is *there*." The crucial question is whether, when *there* refers to a representation that slightly mislocalizes the object in space, due to (e.g.) foveal bias, we still think that S has knowledge about the location of the object, knowledge that we might express either in the form of knowledge-that ("S knows that the object is *there*") or knowledge-wh ("S knows where the object is"). While I maintain that even intuitively we still want to call this knowledge, the more important question for the present argument is whether such cases would pose a skeptical problem if we didn't consider them knowledge. Given that our putative knowledge about the locations of the objects around us constitutes an entirely non-trivial proportion of our total body of putative knowledge—relied upon heavily in navigating our environments and latently generated on a continual basis even when we aren't accessing it—we can understand that a skeptical problem would emerge were we to say that at least much of this isn't actually knowledge. However, factivity does precisely this. As *there* in many cases will refer to a location in space where the object is not, "the object is *there*" is not strictly true, regardless of how close to truth it might be. Thus, we can only retain broad knowledge of the locations of objects in our visual environments if we replace factivity with a weaker constraint on knowledge. In short, while certainly more involved than the argument from naïve approximation, we

⁴ Tempting though it may be to jump directly to belief content, by first specifying the behavior on the basis of which we might attribute such belief content, we might avoid worries about whether the beliefs in question really are slightly false.

⁵ The key exception would be when objects are located in the center of vision and therefore not subject to such mislocalization. While I'm not in a position to address the proportion of cases in which this exception holds, it should be clear that there are still a great many cases in which we perform motor actions and/or think about our visual environments relying on representations that aren't exclusively localized to foveal space. Indeed, there is even evidence that specific types of object representations requiring "large-scale feature integration," like those of buildings, rely heavily on information from outside the fovea (Hasson et al., 2002, 479). This is all that is required to make the sort of skeptical argument we're interested in here.

⁶ Note that Bricker also discusses orientation and size, but here I'll stick to location as, again, I think it's the most convincing and straightforward case.

can understand how these cases of cognitive approximation ultimately motivate the same sort of anti-skeptical argument against factivity.

At this point, one might feasibly raise any number of objections to this argument (see especially Shaffer forthcoming), and while both Buckwalter and Turri (2020a, 95–97) and Bricker (2018, ch. 3) address a number of potential concerns, here I won't attempt a comprehensive response to potential objections. Instead, I simply want to observe that these cognitive approximations allow us to avoid the grain objection left open by Buckwalter and Turri. Recall that this objection worries that the representational content in cases of naïve approximation may not actually be sufficiently fine-grained to be false and that Bricker had even suggested something along these lines. As Buckwalter and Turri acknowledge, the cases they utilize are open to this objection because they lack the empirical evidence necessary to establish the precision of the representational content in these cases. As the instances of cognitive approximation of interest here certainly don't suffer from this empirical shortcoming, we are now in the position to ask whether representational content subjected to foveal bias is sufficiently fine-grained to clearly count as false.

Put simply, the empirical evidence documenting foveal bias attests quite unequivocally that the content of beliefs regarding spatial position is indeed fine-grained enough to be slightly false, not simply approximate. In order to understand why, it is important to consider the precision of the behavior we wish this representational content to explain. Let's take as an example the misreporting of target stimulus location—on the order of magnitude of single degrees of arc in the visual field—observed by Sheth and Shimojo (2001). The best explanation for this behavior is simply that participants consistently misrepresented—albeit slightly—the stimulus location, and it's unclear how we might straightforwardly explain this behavior if we assume representational content so imprecise that it includes the stimulus location. Understanding the purpose of positing such content to be none other than the explanation of this behavior, as is common (see e.g. Nagel, 2013), it is clear that then we want to say that this content in fact is precise enough to misrepresent stimulus location. As this is the very content associated with beliefs and knowledge about spatial location, we can conclude quite confidently that such beliefs are indeed fine-grained enough to be false, with any corresponding knowledge therefore being inconsistent with factivity. Again, this follows directly from the precision of the behavior driven by this content. In this manner, we can understand how empirical cases like foveal bias can provide the tools necessary to address worries about precision and grain.

To conclude this section, I would again reiterate that dismissing this single worry is of course not on its own nearly enough to constitute a comprehensive argument against factivity. Although I find something like the cases of cognitive approximation that result from foveal bias to be a bit more convincing than cases of naïve approximation, I am under no illusions about the limited acceptance an argument like this is likely to find. Beyond more general worries about dropping factivity, and those other objections alluded to above, one might also worry, for example, that naïve approximations and/or beliefs about spatial location don't quite make up the proportion of putative knowledge requisite for generating a severe skeptical problem. While all these worries will of course need to be systematically addressed by the non-factive project, this, again, has not been my aim here. Instead, the purpose

of this section was to strengthen the argument from skepticism in one very specific way—by observing how cases like those of cognitive approximation can solve a significant objection left open by naïve approximations. Now, in the next section, I'll turn to a second argument against factivity, demonstrating now how Buckwalter and Turri strengthen a shortcoming present in Bricker's account.

3 Inaccessibility and the Argument from Intuitive Judgements about Knowledge

In the previous section, we observed that Bricker's empirically grounded argument from cognitive approximation provides us with the tools required to address a major objection left open by the Buckwalter and Turri argument from naïve approximation. Now, in this section, we might observe something of the converse—Buckwalter and Turri also provide us with the tools required to address a key shortcoming of Bricker's account. The shortcoming here is a speculative characterization of our intuitive judgements about knowledge, which not only rests entirely on armchair introspection, but commits to the thesis that the non-factive project is restricted in scope to reflectively inaccessible belief content. However, experimental findings from Buckwalter and Turri (2020b) address both of these shortcomings. Not only do these findings strengthen armchair speculation regarding willingness to attribute knowledge that is inconsistent with factivity, but they importantly demonstrate this using reflectively *accessible* content. Accordingly, these findings both generally support the core assertion of the non-factive project and suggest the potential for a more complex relationship between factivity-violating knowledge⁷ and reflective access than has previously been considered.

Let's begin with Bricker's characterization of the relevant intuitive judgements about knowledge, in which, for certain slightly false beliefs, like in the foveal bias example discussed above, "we still judge that these beliefs are known" (2018, 103). Crucially, however, Bricker thinks that we only ever make judgements like these, inconsistent with factivity, when the misrepresentative content is reflectively inaccessible (see 2018, 50 & 74).⁸ Let's refer to this as the inaccessibility thesis—while there may be some cases in which we might judge that a false belief constitutes knowledge, this will only be when the false content of this belief is reflectively inaccessible. If misrepresentative content is accessible upon reflection, e.g. in the case of measurement, then we don't attribute knowledge for otherwise comparable, slightly false beliefs. For example, if S believes that object O is *there*, with *there* referring to a spatial position one millimeter removed from O's actual position due to foveal bias, we might judge that S still knows that O is *there*/S knows where O is. However, if S forms the same slightly false belief about the location of O, but this time on

⁷ I'll use the term "factivity-violating knowledge" to refer to tokens of knowledge that are inconsistent with knowledge being factive.

⁸ To be clear, this doesn't mean just access to the fact that content is erroneous, but access to the content itself.

the basis of a reflectively accessible measurement, so that S can now report something like, “O is 13 cm from the edge of the table” when in fact O is 13.1 cm from center, the inaccessibility thesis maintains that we no longer judge that S knows where O is.⁹ Bricker even goes so far as to frame the comparative accessibility of the content in the foveal bias case—relatively accessible in comparison to the visuomotor content he discusses first—as a liability for its convincingness, regretting that “we cannot appeal to the total inaccessibility of this content to bolster the argument’s plausibility” (2018, 103). In short, Bricker characterizes our judgements about knowledge as being sometimes inconsistent with factivity, but only when the approximately-true-but-false content isn’t accessible upon reflection.

Unfortunately, as these claims rely exclusively on the individual introspection of a single philosopher, they can’t themselves constitute an especially strong foundation for the non-factive project. Without any appeal to empirical evidence regarding such judgements, studies from experimental philosophy, or agreement from other philosophers, they might easily be dismissed as outliers. This is especially problematic if we want to motivate the non-factive project, as Bricker does, with the argument that “[n]on-factive analyses of knowledge more successfully capture our intuitive judgements than do factive analyses” (2018, 3). Fortunately, however, this is by no means an intractable shortcoming, as Buckwalter and Turri have provided us with the very resources required to address it.

First, on a preliminary basis, we might note that Buckwalter and Turri themselves arguing against the factivity of knowledge constitutes its own kind of confirmation for the idea that we sometimes might judge that false beliefs are known, at least to the extent that it demonstrates that other epistemologists do indeed share in similar intuitive judgements. However, perhaps more important for our purposes are the experimental findings provided by Buckwalter and Turri, which report on lay judgements about cases involving slightly false content (2020b). In this study, participants were provided with a series of short vignettes, which described either true or false-but-approximately-true content, and then asked (among other things) whether such content constituted knowledge. For example, one vignette read the following [variations indicated in brackets]:

Drew is a civil engineer designing a structure. To fit, he must use the length of an existing foundation, to [the exact inch/within 10 inches] of its total length. After thinking carefully, Drew uses [9905/9910] inches. The actual length of the foundation was 9910 inches. (2020b, 3)

Unsurprisingly, participants indicated a much stronger agreement with knowledge attributions—e.g., “Drew knows the length of the foundation”—in the case that the content was strictly true. Crucially, however, in the case that the content was only approximately true, their agreement with knowledge attribution significantly

⁹ In fairness, Bricker does provide a plausible argument that S does something mistaken if she forms a belief on the basis of measurement whose content is comparable to that offered by visual perception (i.e. sufficiently fine-grained to be slightly false; 2018, ch. 4 §1.3). However, as this argument is offered as an explanation of what Bricker first posits are our intuitive judgements, it isn’t directly relevant here.

outstripped their agreement with the assertion that the content was true, with a clear tendency towards agreement with certain non-factive knowledge attributions. To be clear, this data is by no means conclusive. As I will discuss below, I think we might especially worry about the sensitivity of certain judgements about *truth* to pragmatic factors, and clearly more studies investigating factivity are needed. Nevertheless, these findings suggest a willingness to attribute knowledge for some content that isn't strictly true. Accordingly, these findings lend further credence to the proposal that we sometimes judge that false beliefs can nonetheless be known and might argue that knowledge isn't factive on this basis.

Moreover, it is important to observe that the vignettes used by Buckwalter and Turri all involve slightly false content that is accessible upon reflection. For example, in the above "Drew" vignette, Drew can access whether he represents the length as being 9905 or 9910 inches. This offers parallel suggestions for understanding the role of reflective access for our judgments about knowledge in these cases of approximate truth. First, this suggests that the inaccessibility thesis is likely too strong in its proposal that S's reflective access to slightly false content categorically precludes the willingness to attribute knowledge to S—regardless of whether this is taken to include philosophers or laypeople. Not only do Buckwalter and Turri themselves maintain that such cases can still constitute knowledge (2020a), albeit primarily on an anti-skeptical basis, but their experimental results also indicate some corresponding willingness from laypeople. Additionally, however, the relative weakness of the results provided by Buckwalter and Turri raises the intriguing possibility that the inaccessibility thesis wasn't entirely off the mark—it could still be that accessibility plays a role in these judgements. Or, perhaps we will ultimately want to commit to something like the inaccessibility thesis—not as a descriptive claim about judgements about knowledge but instead a principle about which false beliefs can be known—to avoid certain new arguments levied by Shaffer against the non-factive project (forthcoming §§3 & 5). While unfortunately beyond the scope of the present paper, it is worth keeping such an option in mind going forward. In this way, we might understand how the inaccessibility thesis may yet provide an important theoretical distinction in the continued development of the non-factive project, highlighting the need for future study.

In the end, although again not a knock-down argument against factivity, we might understand how experimental results from Buckwalter and Turri lend additional credence to a central component of the non-factive project. Moreover, while these results suggest that the inaccessibility thesis is likely too strong—at least as a descriptive claim about the judgements of lay people—it nevertheless does raise important points about the relationship between accessibility and factivity-violating knowledge.

4 Replacing Factivity: Truthlikeness vs. Representational Adequacy

In the previous two sections, we reviewed variations of two key arguments against factivity offered by Bricker and Buckwalter/Turri, observing how they mutually reinforce one another. In this way, we can understand how facilitating interaction

between these accounts serves to strengthen the non-factive project as a whole. Now, in this section, I want to further advance the non-factive project by addressing the question of how it ought to proceed with replacing the factivity constraint on knowledge. I'll first consider the truthlikeness condition proposed by Bricker and then the representational adequacy account proposed by Buckwalter and Turri, after which I'll argue that truthlikeness displays a key advantage over representational adequacy—while the adequacy approach struggles to preserve the close connection between knowledge and truth, this isn't a problem for truthlikeness. As the truth-knowledge connection is widely regarded as fundamental, including by Buckwalter and Turri, this suggests that truthlikeness represents the more promising starting point for the future development of the non-factive project, as it seeks to solidify a replacement condition for factivity. Here, I will refrain from committing to any definitive conclusions about whether truthlikeness itself might constitute a complete replacement for factivity, or instead if some additional elements (perhaps even a version of representational adequacy) might also be needed.

Let's begin with the truthlikeness account suggested by Bricker, which proposes that we replace factivity with a *truthlikeness* condition on knowledge (2018, 72–74). The idea here is itself quite simple. Taking inspiration from Popper (1963), here truthlikeness is just a scalar measure of the degree to which propositional content is true. That is, truthlikeness comes in degrees, with (e.g.) beliefs closer to truth being more truthlike, strictly true beliefs being maximally truthlike, and so forth. Accordingly, this account maintains that rather than requiring strict truth, knowledge instead only requires some sufficient degree of truthlikeness. As with other such concepts that come in degrees—like justification, modal closeness, or, as we'll discuss shortly, representational adequacy—it isn't entirely clear how or where to set the threshold for truthlikeness sufficient to constitute knowledge. However, Bricker maintains that “[i]n order for this to be a plausible replacement for factivity, the threshold for sufficient truthlikeness needs to be very high” (2018, 73). As we will see below, while one might appeal to the requirements of our cognitive systems in explicating just how high the truthlikeness threshold for knowledge might be, this approach ultimately takes truthlikeness itself to be the central concept governing whether beliefs stand in the relation to truth requisite of knowledge.

In contrast with the truthlikeness account, Buckwalter and Turri propose that we substitute *representational adequacy* for the factivity constraint on knowledge (2020a, 97; 2020b, 2). Although they don't offer an excess of detail in explaining precisely what adequacy means here, they do present one key way in which a representation might adequately represent the truth—being able to successfully guide goal-oriented actions:

[R]epresentations need not be true in order to count as knowledge. Instead, they only need to *adequately represent* the truth. Although there are potentially many ways that approximations could be adequate, one way is for them to serve our purposes well enough to facilitate action and help us to achieve our goals in a particular circumstance. (2020a, 97; original emphasis)

Buckwalter and Turri continue by illustrating this idea with the example of approximating the number pi:

For example, 3.14 might be adequate, and hence known, as the value of pi in the grade school classroom but inadequate, and hence not known, as the value of pi in the lab engineering a global positioning system. (2020a, 97)

While I'm not sure whether I'm convinced that a school child can actually know that pi is 3.14, this is a bit beside the point. The main idea here is just that while the version of the account developed by Buckwalter and Turri leaves open the possibility of other routes to representational adequacy, this pragmatic adequacy for goal-oriented action guidance represents the way in which they explicitly formulate it.

While this conclusion that pragmatic adequacy is a sufficient replacement for factivity differs considerably from the truthlikeness approach,¹⁰ there is an element of overlap between the two accounts worth discussing here. Not only are they not mutually exclusive, but a sense of representational adequacy might even be incorporated into specific versions of a truthlikeness condition. Notably, one way to understand the degree of truthlikeness required for (at least) low-level perceptual beliefs to be known is via the output requirements of perceptual-belief-generating processes, e.g., the requirement that output beliefs can successfully guide motor action (Bricker, 2018, 80–83). Contrasting with the formulation of representational adequacy favored by Buckwalter and Turri, however, this approach considers success conditions at the systems level, not the level of individual representations. What it means for a belief to—borrowing the language of Buckwalter and Turri—adequately represent the truth is not simply a matter of being true enough for the goals associated with an individual belief. As we'll see shortly, this can sometimes be woefully insufficient, as certain individual goals have remarkably low truth requirements. Instead, on this approach to truthlikeness, beliefs need to fulfill the truth requirements of their respective belief-forming systems. While it is a subtle distinction, this emphasis on the system level vs. the representational level may at least partially explain the fundamental structural difference between truthlikeness and representational adequacy accounts: Truthlikeness is easier to understand as a systems-level requirement than pragmatic adequacy, and, conversely, pragmatic adequacy makes more sense when understood at the level of individual representations.

As it turns out, the systems-level truthlikeness account avoids a key problem we might identify for the representation-level pragmatic adequacy account—truthlikeness preserves the vital truth-knowledge connection in a way that representational adequacy does not. As both Bricker (2018, 74) and Buckwalter and Turri (2020a, 97) note, there is a strong intuitive connection between knowledge and truth, which both accounts make a point of seeking to retain. That non-factive theories of knowledge would still wish to acknowledge that knowledge is very closely related to truth

¹⁰ To be clear, Buckwalter and Turri leave the door open to the possibility that other types of adequacy might sometimes be required to explain knowledge, and in this sense, they do not contend that this pragmatic adequacy is a sufficient replacement for factivity. However, in framing representational adequacy as a sufficient replacement for factivity, and pragmatic adequacy as one type of representational adequacy, the implication is that when false beliefs display this pragmatic adequacy, they constitute knowledge to the same extent as otherwise comparable true beliefs. This is what I mean by “a sufficient replacement for factivity.”

should come as no surprise. After all, at least traditionally, all those paradigmatic instances of known belief we might identify are cases of true belief. Moreover, critically, all those examples of potentially known false belief identified above are cases in which the beliefs are still very close to being true. All of this gives us good reason to think that a very specific sort of close connection with truth is an integral feature of knowledge—a token-level connection, on which no instance of knowledge can fail to be in some way closely related to the truth. Accordingly, it is critical that in replacing factivity with some weaker constraint, we don't produce an account on which beliefs that are wildly decoupled from the truth might qualify as knowledge.

As it happens, however, the representational adequacy account does just this. While representational adequacy might be generally associated with truth, this truth connection can suffer a catastrophic breakdown in certain cases. We might illustrate this by first observing that there will often be goals for which representing the truth sufficiently adequately to accomplish that goal is quite a low bar to clear. Consider the following example:

During her inaugural visit to the Cairngorms, Anna's goal is to hike from Blair Atholl to the Falls of Tarf. From her internet research, Anna believes the falls to be located around 60 miles straight up Glen Tilt¹¹ from Blair Atholl. In the course of her research, she also gains a sufficient understanding of the topography of the glen to successfully orient herself and hike up it without easily losing her way, as well as sufficient familiarity with the distinctive appearance of the falls to easily recognize them on sight. Note that prominent signage identifying the falls also serves to ensure that hikers know when they have arrived. Anna sets out from Blair Atholl up Glen Tilt, well outfitted for a cross-country trek into the Scottish Highlands, and easily reaches the falls on the basis of her belief regarding its location, albeit much sooner than she expected. As it happens, the Falls of Tarf are only around 12 miles up Glen Tilt.

An example like this highlights a key characteristic of pragmatic adequacy. While oftentimes a close proximity with truth is required for a belief to successfully guide an agent's actions towards her goal, there is no reason that this needs to be the case. Anna's belief about the location of the Falls of Tarf isn't even close to the truth, but in this case, it is still sufficient to successfully guide her actions. Moreover, this decoupling of representational adequacy from truth explains why, in all likelihood, we don't judge that Anna knows (i) where the Falls of Tarf are or (ii) that the Falls of Tarf are 60 miles up Glen Tilt. As discussed above, even if knowledge isn't strictly factive, it is clear that there is still some close connection between cases of knowledge and the truth. However, here we can understand that representational adequacy and truthlikeness vary significantly in the extent to which they preserve this truth-knowledge connection. Truthlikeness by definition retains the truth-knowledge connection, as any known false beliefs will still be very close to truth. On the other

¹¹ Glen Tilt is a deep, narrow, and impressively straight valley that straddles the River Tilt all the way up to the Falls of Tarf. This specific topography is actually quite important for our example, as it means that once Anna begins her hike up the glen, she is unlikely to become disoriented.

hand, representational adequacy drops the ball in cases where beliefs quite distant from the truth are still able to successfully guide action. Replacing factivity with representational adequacy alone would therefore result in a theory of knowledge that struggles to preserve the truth-knowledge connection. Now, at this point, I cannot rule out the possibility that some additional condition might be added to the representational adequacy account to facilitate the same theoretical work of connecting knowledge and truth traditionally underwritten by factivity, but this is precisely the point. Representational adequacy, in contrast with truthlikeness, is clearly unable to itself do the theoretical work requisite of a replacement for factivity. For this reason, I maintain that truthlikeness, not representational adequacy, represents the more promising starting point as we move forward with developing a replacement for factivity.

Before concluding this section, I want to quickly respond to three worries that advocates of the representational adequacy account are likely to raise in response to my argument. First, it is likely that those of a more hardened experimentalist bent will not be especially amenable to my more armchair approach, especially where I refer to putative judgements about knowledge without any experimental evidence. For this reason, I have been careful to not use any such judgements as a core component of my argument. While I expect that many more traditional epistemologists are likely to agree that clearly Anna doesn't have knowledge, and therefore be skeptical of the representational adequacy account on this basis, it is not required that you do so for my argument to go through.¹² Instead, all that is needed is that we can observe the breakdown of the truth-knowledge connection when representational adequacy is applied to the case, which might be done without committing to any armchair judgements about whether the particular case actually constitutes knowledge.

Next, one might object to the characterization of the representational content that I have identified as guiding Anna successfully to her goal. While I have specifically maintained that her actions are guided by her representation of the Falls of Tarf as being 60 miles up the glen, one might instead contend that it is simply her representation of the falls as up Glen Tilt that guided her successfully to them. As her belief in this case would not just be close to truth but strictly true, the objection goes that I have not in fact described a case in which representational adequacy and truth decouple. However, the problem with this objection is that her belief regarding the distance of the falls from Blair Atholl, while wildly inaccurate, is absolutely crucial in facilitating her success in reaching her terminus. Had she instead believed that the falls were only 500 yards up the glen, or instead had no real belief about their distance from Blair Atholl at all, it is highly unlikely that she would have been either physically or mentally prepared for a 12-mile trek through increasingly unforgiving

¹² Although even then things aren't so cut and dry. The representational adequacy account is of course not committed to the idea that all adequate representations must be known, so one would need to specify non-factive versions of familiar constraints on knowledge to evaluate whether beliefs in cases like Anna's might fail to constitute knowledge exclusively in virtue of their disconnect with the truth. I am thankful to an anonymous reviewer for raising this point.

highland terrain, and she almost certainly wouldn't have brought the supplies required for both the outgoing and return journeys. Quite clearly, Anna's belief that the falls were up the glen is not sufficient to explain how she was able to successfully reach them. Instead, we must acknowledge that her belief that they were 60 miles from Blair Atholl played a crucial role in her physical and mental preparation, determining the supplies that she took with her, and ultimately her success.

Finally, moving on from this specific example, one might also object that the experimental findings from Buckwalter and Turri discussed above provide evidence that representational adequacy, not truthlikeness, is indeed the operative concept employed by laypeople in their judgements about knowledge. After all, the primary claim of the paper reporting those results is that these findings support the representational adequacy account, which follows from their observation that participants agreed more strongly with non-factive knowledge attributions when the content of those attributions was adequate for some pragmatic purposes (2020b). While this objection is understandable, I would respond by calling attention to the following: In that study, for cases in which the content depicted in vignettes was clearly and explicitly false, participants also more strongly agreed that this content was *true* in the cases of pragmatic adequacy when compared to cases of inadequacy. Moreover, although Buckwalter and Turri didn't analyze the data in order to address this specific issue, for cases in which content was *false*, visual inspection of their results indicates a correlation between (i) agreement that the content was true and (ii) agreement that the content constituted knowledge. This suggests that there are two issues with participant responses that preclude the straightforward conclusion that this study supports representational adequacy over truthlikeness. First, when content was false, participants did not strongly disagree with the statement that the content was true with the consistency we might expect from reliable evaluators. Moreover, relatedly, this seeming correlation of agreements with statements of truth and attributions of knowledge indicates that truth may actually be playing a more important role than Buckwalter and Turri are acknowledging. As a matter of speculation, if we might reasonably understand participants slightly or somewhat agreeing with a statement of truth as indicating something like an evaluation of truthlikeness, this apparent correlation would then also constitute evidence for the truthlikeness account. While this is of course a matter of future research, not something I can definitively settle here, that really isn't the point. Instead, the point is simply that while the findings from Buckwalter and Turri may provide some evidence for representational adequacy over *factivity*, they do not obviously do the same for representational adequacy over truthlikeness.

To conclude, I'd simply like to reiterate that my intention with this section was simply to take a step in the direction of establishing how the non-factive project ought to go about replacing the factivity constraint, not to definitively settle the matter. As I have argued here, representational adequacy is likely a poor starting point, given its weak and entirely too contingent connection with truth. Moreover, while truthlikeness maintains a much stronger, necessary relationship with truth, allowing us to preserve the critical truth-knowledge connection, this of course does nothing like entail that truthlikeness is conclusively the singular best option for replacing the

factivity constraint—let alone that knowledge actually is factivoid.¹³ Additionally, it could still be that representational adequacy has a role to play in the replacement of factivity, perhaps in some sense fixing the degree of truthlikeness required for knowledge (discussed above). Nevertheless, I think we might at least at this point reasonably conclude that truthlikeness is the most promising extant option for future development available to the nascent non-factive project.

5 Conclusion: Moving the Non-factive Project Forward

The aim of this paper was to advance the non-factive project by facilitating the interaction between recent clusters of arguments and accounts all dedicated to pushing back against the factivity condition on knowledge (Bricker, 2018; Buckwalter & Turri, 2020a, b). Ultimately, we observed that (i) cognitive approximation offers a solution to the grain problem present in the argument from skepticism derived from naïve approximation; (ii) empirical findings from Buckwalter and Turri suggest that the inaccessibility thesis is likely mistaken; and (iii) in preserving the fundamental connection between knowledge and truth, the truthlikeness framework avoids a significant problem associated with the representational adequacy account. All told, this provides an encouraging assessment of the present state of the fledgling non-factive project, an epistemological research program that, just a few years ago, would have been difficult to foresee.

Nevertheless, as I have been careful to repeat throughout this paper, nothing I have said here constitutes a comprehensive, airtight argument against the factivity of knowledge. Indeed, here I have identified at least four open issues for the non-factive project to address. First, there is the general experimental question concerning the strength and ubiquity of lay agreements with non-factive knowledge attributions, the present evidence for which is admittedly a bit soft. Second, there are specific experimental questions regarding whether, for cases of slightly false belief, (i) reflective access modulates knowledge attributions or (ii) the increased agreement with statements of truth reflects an assessment of truthlikeness, or instead is evidence of unreliable evaluation. Next, there is the trickier, more conceptual question of the proportion of our putative knowledge that is actually slightly false. While this is a crucial element of both the naïve and cognitive approximation versions of the argument from skepticism, I'm not sure that this has been established as strongly as it might be. Finally, there is the much broader question of just how much theoretical work is performed by the factivity constraint and whether this work might then be satisfactorily achieved by something weaker. While here I specifically focused on the connection to truth facilitated by factivity (§4), this almost certainly does not exhaust the theoretical work carried out by factivity for theories of knowledge. In the end, while everything surveyed here gives us a good reason to be optimistic about the future

¹³ We might use “factivoid” to refer to non-factive accounts of knowledge on which factivity is replaced with a strict truthlikeness condition (see Bricker 2018, 72).

prospects of the non-factive project meeting such challenges, it is clear that most of its work still lays ahead.

Funding This work was supported by Sven Bernecker's Alexander von Humboldt Professor grant. Open Access funding enabled and organized by Projekt DEAL.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Ackermann, R. (1972). *Belief and knowledge*. Macmillan.
- Annis, D. (1974). Belief and knowledge. *The Philosophical Quarterly*, 24(94), 81–82.
- Armstrong, D. (1973). *Belief, truth and knowledge*. Cambridge University Press.
- Audi, R. (2010). *Epistemology*. Routledge.
- Bocianski, D., Müsseler, J., & Erlhagen, W. (2008). Relative mislocalization of successively presented stimuli. *Vision Research*, 48(21), 2204–2212.
- BonJour, L. (2010). *Epistemology classic problems and contemporary responses* (2nd ed.). Rowman & Littlefield Publishers Inc.
- Brenner, E., Mamassian, P., & Smeets, J. (2008). If i saw it, it probably wasn't far from where i was looking. *Journal of Vision*, 8(2), 7.1-10.
- Bricker, A. (2018). *Visuomotor noise and the non-factive analysis of knowledge*. PhD Thesis. The University of Edinburgh.
- Bricker, A. (2019). Postnarrativism, Historiographical Evaluation, and Truth. *Journal of the Philosophy of History*, 1–19.
- Buckwalter, W. (2014). Factive verbs and protagonist projection. *Episteme*, 11(4), 391–409.
- Buckwalter, W., & Turri, J. (2020a). Knowledge and truth: A skeptical challenge. *Pacific Philosophical Quarterly*, 101(1), 93–101.
- Buckwalter, W., & Turri, J. (2020b). Knowledge, adequacy, and approximate truth. *Consciousness and Cognition*, 83, 102950.
- Davidson, D. (1988). *Epistemology and Truth. Subjective, Intersubjective, Objective*. Clarendon Press.
- Fortenbaugh, F., & Robertson, L. (2011). When here becomes there: Attentional distribution modulates foveal bias in peripheral localization. *Attention, Perception, & Psychophysics*, 73(3), 809–828.
- Hannon, M. (2013). 'Knows' entails truth. *Journal of Philosophical Research*, 38, 349–366.
- Hasson, U., Levy, I., Behrmann, M., Hendler, T., & Malach, R. (2002). Eccentricity bias as an organizing principle for human high-order object areas. *Neuron*, 34(3), 479–490.
- Hazlett, A. (2010). The myth of factive verbs. *Philosophy and Phenomenological Research*, 80(3), 497–522.
- Ichikawa, J. and Steup, M. (2018). The analysis of knowledge. *The Stanford Encyclopedia of Philosophy* (Summer 2018 Edition), Edward N. Zalta (ed.), URL = <<https://plato.stanford.edu/archives/sum2018/entries/knowledge-analysis/>>.
- Kerzel, D. (2002). Memory for the position of stationary objects: Disentangling foveal bias and memory averaging. *Vision Research*, 42(2), 159–167.
- Nagel, J. (2013). Knowledge as a mental state. *Oxford Studies in Epistemology Volume 4*. Oxford University Press.
- Niiniluoto, I. (1999). *Critical scientific realism*. Oxford University Press.

- Nozick, R. (1981). *Philosophical explanations*. Clarendon Press.
- Ono, F., & Watanabe, K. (2011). Attention can retrospectively distort visual space. *Psychological Science*, 22(4), 472–477.
- Odegaard, B., Wozny, D., & Shams, L. (2015). Biases in visual, auditory, and audiovisual perception of space. *PLoS Computational Biology*, 11(12), e1004649.
- Popper, K. (1963). *Conjectures and refutations: the growth of scientific knowledge*. Routledge.
- Pritchard, D. (2017). *Epistemology* (2nd ed.). Palgrave Macmillan Limited.
- Pritchard, D., & Kelp, C. (2010). *Two deflationary approaches to Fitch-style reasoning. New Essays on the Knowability Paradox*. Oxford University Press.
- Shaffer, M. (forthcoming). Can knowledge really be non-factive? *Logos and Episteme*.
- Sheth, B., & Shimojo, S. (2001). Compression of space in visual memory. *Vision Research*, 41(3), 329–341.
- Shope, R. (1983). *The analysis of knowing: A decade of research*. Princeton University Press.
- Suzuki, S., & Cavanagh, P. (1997). Focused attention distorts visual space: An attentional repulsion effect. *Journal of Experimental Psychology: Human Perception and Performance*, 23(2), 443–463.
- Turri, J. (2011). Mythology of the factive. *Logos & Episteme*, 2(1), 141–150.
- Uddin, M., Kawabe, T., & Nakamizo, S. (2005). Attention shift not memory averaging reduces foveal bias. *Vision Research*, 45(25), 3301–3306.
- Zimmermann, E., Fink, G., & Cavanagh, P. (2013). Perifoveal spatial compression. *Journal of Vision*, 13(5), 21–21.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.