Visuomotor Noise and the Non-Factive Analysis of Knowledge

Adam Michael Bricker

A Thesis for the Degree of Ph.D. (Philosophy)
The University of Edinburgh
2018

I, the undersigned, declare that: (i) the thesis has been composed by me, (ii) the work is my own, (iii) the work has not been submitted for any other degree or professional qualification except as specified, and (iv) that any included publications are my own work, except where indicated throughout the thesis and summarised and clearly identified on the declarations page of the thesis.

Adam Bricker

Abstract:

It is all but universally accepted in epistemology that knowledge is factive: S knows that p only if p. The purpose of this thesis is to present an argument against the factivity of knowledge and in doing so develop a non-factive approach to the analysis of knowledge. The argument against factivity presented here rests largely on empirical evidence, especially extant research into visuomotor noise, which suggests that the beliefs that guide everyday motor action are not strictly true. However, as we still want to attribute knowledge on the basis of successful motor action, I argue that the best option is to replace factivity with a weaker constraint on knowledge, one on which certain false beliefs might still be known. In defence of this point, I develop the non-factive analysis of knowledge, which demonstrates that a non-factive constraint might do the same theoretical work as factivity.

Lay Summary:

Do you know where your hands are? Or where the objects around you are? Or how many "7" is? It certainly seems like it, doesn't it? But the way that our brains work means that we need to think about this a bit more. What if things aren't actually where it looks like they are, but they're still very close to where it looks like they are? What if how big we think numbers are is not exactly how big they actually are? We can still say that we know these things, but we might have to think about knowing things differently than we did before.

Table of Contents

Chapter One: Background and Introduction

This chapter has two primary aims. The first of these is to introduce the project of this thesis and provide a general outline of its argument. This will be addressed in section one, in which I'll discuss, chapter-by-chapter, the key arguments of subsequent chapters. In doing so, I hope to make clear both in what sense and on what grounds I advance the non-factive analysis of knowledge. The second aim of this chapter is to provide an overview of relevant background information. This will include extant philosophical discussion of the factivity of knowledge (§2), localizing this project within epistemology, especially with respect to epistemological naturalism (§3), and introducing cognitive neuroscience and its influence on this project (§4).

1. Chapter Overview

The purpose of this thesis is to argue in favour of a non-factive analysis of knowledge. At its most basic, this project might be thought of as reducing to one key claim: Non-factive analyses of knowledge more successfully capture our intuitive judgements than do factive analyses.[1] There are two sorts of ways I'll argue for this claim. First, I'll argue that non-factive analyses of knowledge can account for intuitive judgements that factive analyses cannot account for. That is, in certain cases our intuitive judgements about knowledge conflict with factivity. This is the most demanding aspect of this project, and encompasses material in chapters two, three, and five. Second, I'll also argue that factive analyses of knowledge do not capture any intuitive judgements that a non-factive analysis cannot also capture. This is the primary task of chapter four. In short, we might replace factivity with a weaker, non-factive constraint on knowledge without losing any ability to capture our intuitions. This is what I call the "viability" of the non-factive analysis of knowledge.

Before continuing with an overview of each chapter, I briefly want to note two things. First, the above is of course only a rudimentary sketch and is intended only to play a familiarizing role. Second, my arguments for non-factive knowledge rest largely on empirical evidence. This empirical evidence is used primarily to establish specific features of the representational content that forms the basis of knowledge, as I discuss below.

---

[1] It is worth noting that this isn't the only argument advanced for non-factive knowledge. Developed in parallel with it is a secondary argument that factive analyses invite a skeptical problem that non-factive analyses might easily avoid. This anti-skeptical argument functions as something of a stopgap for the primary argument from intuitive judgements. That is, even if you don't grant there are judgements about knowledge that violate factivity, there remains the problem that factive analyses seem to invite a rather troubling form of skepticism.

1.1 Chapter Two: Empirical Foundations

Chapter two provides a detailed account of the primary empirical evidence on which my subsequent argument against the factivity of knowledge is based. It assumes a simple structure, detailing the empirical support for four key claims:

(1) Perception can occur in the absence of consciousness.

(2) The same information is often represented internally in a multitude of ways.

(3) The mental content that guides motor action is not reflectively accessible.

(4) The mental content that guides motor action is not perfectly accurate.

As I'll discuss in depth, all four of these claims are widely accepted in cognitive neuroscience, and enjoy a broad range of support from behavioural, neuropsychological, and neuroimaging evidence. The first two of these claims play a preliminary, background role. While my subsequent arguments will not directly rely on them, without first establishing them, it's unclear how one might understand the subsequent claims. These latter claims, (3) and (4), provide the empirical foundation for the argument against factivity. Together, they paint a picture of the mental content guiding motor action that appears quite different from what we might otherwise assume when considering knowledge and action. The most important role in my argument is played by claim (4). As I argue in chapter three, this inaccuracy in the mental content guiding motor action, "visuomotor noise," poses an intractable problem for the factivity of knowledge.

1.2 Chapter Three: Against the Factivity of Knowledge

In chapter three I present my primary argument against the factivity of knowledge. The argument is structured as three initially plausible assumptions that lead to an undesirable conclusion:

(1) The mental content that guides motor action is (a) reflectively inaccessible and (b) not perfectly accurate.

(2) Perceptual beliefs guide motor action.

(3) Knowledge is factive.

(C) The everyday perceptual beliefs that guide successful motor action cannot be known.

Assumption (1) is the conjunction of claims (3) and (4) from chapter two. Their initial plausibility derives from the significant empirical support detailed in chapter two. Assumption (2) is initially plausible as it is simply the product of our ordinary practice of belief attribution. Under ordinary circumstances, we readily attribute certain perceptual beliefs on the basis of motor action alone. Finally, as I'll discuss in the next section, the initial plausibility of assumption (3) reduces to it being all but universally accepted in epistemology.

The conclusion (C) is undesirable on multiple fronts. First, since we often judge that the beliefs we attribute on the basis of motor action are indeed known, accepting (C) requires us to go against our intuitive judgements about knowledge. Moreover, accepting (C) is a type of skeptical result, insofar as it means accepting that we have less knowledge than we ordinarily assume we do. Over the course of chapter three, I argue at length, not only that (C) follows from (1)-(3), but also that the best option is rejecting assumption (3). That is, this is a real problem, and the best solution to this problem is to reject the factivity of knowledge.

Rather than discussing the details of this argument, here I simply want to highlight the key ways in which it supports the overall project of this thesis. While one might simply grant that of course an argument against the factivity of knowledge motivates the non-factive analysis of knowledge, the specifics are a bit more complex than this. To see this, it is important to distinguish between three claims:

(i) Factive accounts of knowledge invite a skeptical problem that non-factive accounts don't.

(ii) Factive accounts of knowledge cannot explain judgements about knowledge that non-factive accounts can.

(iii) We should reject the factivity of knowledge.

Chapter three will make all these claims, specifically using (i) and (ii) in support for (iii). However, it is important that we don't take (iii), the conclusion of the chapter three, to itself support the argument in subsequent chapters. Rather, it is (i) and (ii) that play a direct role in the argument for a non-factive analysis of knowledge. In fact, in this project, the fact that a non-factive analysis of knowledge works better than a factive one is taken to support (iii), not the other way around.

1.3 Chapter Four: The Factivoid Analysis of Knowledge

Chapter four is devoted to the development of the non-factive analysis of knowledge. The primary argument I advance in this chapter is that the non-factive analysis of knowledge does not suffer from any novel misattribution of knowledge. That is, non-factive analyses of knowledge can explain all the same intuitive judgements that factive analyses do. As already mentioned, I call this the viability of the non-factive analysis of knowledge. In order to demonstrate this, I present three candidate non-factive analyses of knowledge. The first of these is a non-factive version of justified true belief (JTB), and the second is a non-factive version of robust virtue epistemology (RVE). While neither of these non-factive proposals solves the problems of the factive analyses on which they are based, the important point here is that they don't encounter any novel misattribution. In short, in addition to handling cases, like those involving visuomotor noise, that factive analyses cannot, non-factive analyses can also handle all the cases their factive counterparts can. This is the central idea behind this chapter, and indeed this project.

Additionally, I also propose a third, especially speculative account of knowledge that, while utilizing the RVE framework, represents a radical departure from standard epistemology. This account seeks to unify propositional knowledge, practical knowledge (knowledge-how), and psycho-cognitive knowledge (the sense in which "knowledge" is used in large bodies of empirical literature) all under a single concept. The key point here is that this sort of unifying project is something with obvious advantages that are only attainable under a non-factive conception of knowledge. My intention is by no means to argue that this is a perfect analysis of knowledge. Indeed, my design with this chapter is not to advance any account with the intention that we take it to be ***the*** analysis of knowledge.[2] The purpose of this chapter is to argue that if there is such a thing as ***the*** analysis of knowledge, it is going to be a non-factive.

1.4 Chapter Five: Generalization

The final chapter of this thesis covers topics that might be collectively labelled "generalization." The discussions contained within this chapter are largely independent from one another. Most of the chapter involves additional empirical arguments against the factivity of knowledge. Looking beyond visuomotor noise, we find that other bodies of empirical research suggest a similar sort of argument. The most promising of these emerges from the noisy analogue magnitude representations that guide judgements about numerical and temporal magnitude. I also discuss examples from proprioception and other cases of perception. Here it's important to note that the promise of additional arguments against factivity plays a key role in the argument presented in previous chapters. The larger the set apparent knowledge that cannot be captured by a factive analysis, the more severe the skeptical problem invited by retaining factivity. Since my argument against accepting the factivity of knowledge relies in part on our desire to avoid any serious risk of skepticism, this argument is strengthened by this chapter's contribution on new empirical problems for factivity. Additionally, I end the chapter by briefly addressing future prospects for further generalization.

2. The Factivity of Knowledge

It is all but universally accepted in epistemology that knowledge is factive: S knows that p only if p. As put by Armstrong:

> I do not think that there has ever been any serious doubt that it is a necessary condition for knowledge that what is known be true. (1973, 137)

---

[2] We might of course note that there is no widely accepted analysis of knowledge, and all prominent proposals seem to come with notable shortcomings (see Ichikawa & Steup 2018).

As put by Davidson:

> Everyone agrees that what is known must be true. (1988, 177)

As put by Turri:

> One of the very few things that epistemologists of all stripes agree on is that knowledge requires truth – or as it's often put, that knowledge is factive. We have seen controversy over whether knowledge requires justification and whether knowledge requires belief, but not whether it requires truth. (2011, 141)

As put by Pritchard and Kelp:

> Factivity seems to play an important—indeed, indispensable—role in any plausible theory of knowledge. (2010, 330)

As put by Plato:

> ***Socrates:*** And will a man ever have knowledge of anything the truth of which he fails to attain?
> ***Theaetetus:*** How can he, Socrates? (186c-d)

In short, the idea that knowledge is factive dates as far back as Plato[3] and remains epistemological orthodoxy to this day. Beyond this, there is little more to say about the fact that almost everyone in philosophy accepts that knowledge is factive, and I won't belabour the point more than I already have. Instead, in this section I'll provide a brief overview of the standard arguments taken to motivate factivity, after which I'll discuss two notable examples in which the orthodoxy has, to varying degrees, been challenged: Ackerman's "Pragmatic Analysis" (1973) and Hazlett's "Myth of Factive Verbs" (2010). In doing so, I hope to both introduce the main challenges a non-factive analysis of knowledge might face, as well as distinguish the claims of this project from those made previously.

2.1 Why Factivity?

The most frequently cited motivation for the factivity of knowledge is that it is able to explain our intuition that utterances both like and of the form "S knows that p, and not p" seem incoherent or contradictory (Ackermann 1973; Kalpan 2006; Tsohatzidis 2012; Hannon 2013). If knowledge is factive, then S knowing that p entails p. Thus, such utterances seem contradictory simply because they are contradictory. While I will return to this argument shortly (§2.3), as well as in chapter four, section four, there is little more to say about it here. In the light of the arguments that will unfold in the interim, upon returning to this matter I hope to have demonstrated that this doesn't offer a compelling argument in favour of factivity.

[3] See specifically the Meno and the Theaetetus.

Although it's easy to confuse with the above argument, there is a second, distinct argument one might give in defence of the factivity of knowledge. Rather than appealing to our linguistic judgements about when utterances seem contradictory, one might instead appeal to our epistemic judgements about when beliefs are known. Upon surveying all those beliefs we take to be known, it doesn't appear that any known beliefs are false. If we take a known belief and stipulate that it's false, we no longer judge that it is known. For example, if we begin with my knowledge that the capital of Scotland is Edinburgh, and then imagine a counterfactual in which the capital is secretly moved to Aberdeen, we no longer judge that I know that the Edinburgh is the capital. Conversely, if we take a seemingly false belief that otherwise seems to fit the bill for knowledge, and stipulate that it is true because of this, we then judge that that belief is known.

All these observations about our epistemic judgements can be explained if we take knowledge to be factive, and it's unclear how we might otherwise explain them without factivity. In short, factivity plays a significant explanatory role for our theories of knowledge. This is what I think Pritchard and Kelp are getting at when they talk about the "important—indeed, indispensable—role in any plausible theory of knowledge" played by factivity (2010, 330). Moreover, at least from an epistemological perspective, I think this is the primary challenge to any non-factive account of knowledge: It needs to replicate the explanatory power of factivity by different means. Accordingly, this is the primary challenge my non-factive project is geared to face. Especially in chapter four, I'll argue that non-factive accounts of knowledge can be just as explanatorily successful as factive ones, indeed more so.

Finally, in the support of the factivity of knowledge one might also appeal to a general intuitive sense that knowledge is factive. That is, knowledge just seems like a factive sort of thing. While I won't dispute that this general intuition is likely widespread, we might note here that general intuition is not all that evidentially valuable. It famously might be overridden by both specific epistemic judgements, as in Gettier cases (1963), and for anti-skeptical reasons, as in the case of infallibilism (see Cohen 1988). As my argument against factivity appeals both to epistemic judgements and anti-skepticism, we can consider it to simply override any general intuitive sense we might otherwise have about knowledge being factive.

2.2 Ackermann's Pragmatic Analysis

As far as I can determine, a genuinely non-factive analysis of knowledge has been proposed exactly once in contemporary epistemology. This is the so-named "Pragmatic Analysis" offered by Robert J. Ackermann in his obscure 1973 monograph *Belief and Knowledge*. Apart from a perplexed reviewer remarking, "He seems to attack the truth condition . . . It is not clear to me why he does this." (Annis 1974, 82), Ackermann's analysis doesn't appear to have attracted any significant attention, either contemporarily or in the decades since. As I will discuss in this section, there is likely good reason for this. To be clear, this thesis is in no way inspired by Ackermann's analysis, nor does it share any meaningful similarities beyond being non-factive. In fact, I maintain that it appears that Ackermann's argument is obviously flawed. Nevertheless, given that it is the rare example of a non-factive analysis of knowledge in the wild, I think it is worth discussing here, if only for historical purposes. At the very least, it provides a good example of the pitfalls such attempts might stumble into.

In *Belief and Knowledge*, Ackermann distinguishes between two different types of analyses of knowledge. The first type of analysis of knowledge is the familiar one on which knowledge entails belief and truth, which he dubs the "Ideal Analysis" (1973, 74). However, after introducing the Ideal Analysis, he maintains, "There is a line of attack on [the factivity constraint] that has not received much attention in the literature" (ibid, 74).[4] This is the argument that ultimately leads him to a second type of analysis, the non-factive "Pragmatic Analysis" of knowledge (ibid, 76). The starting point for this argument is a "partial analysis" he provides for first-person knowledge assertions:

> To say "I know that p" is to assert that one can meet all relevant non-metaphysical objections to p. (ibid, 67)

Although Ackermann (somewhat perplexingly) doesn't refer to it as such, this is simply a relevant alternatives account of knowledge, on which "knowing *p* requires ruling out all *relevant* alternatives to *p*" (Luper 2016). Most noteworthy about this claim is that it has a strong intuitive appeal (see Carter and Pritchard 2016, §1) but famously runs into problems when coupled with epistemic closure (Dretske 1970). Especially considering that Ackermann fails to even acknowledge these problems, we might note this to be the first major shortcoming of the Pragmatic Analysis.

Ackermann then notes that on the above partial analysis, saying "I know that p" can mean asserting either of the following (emphasis mine):

> One can meet all ***current*** relevant non-metaphysical objections to *p*.
>
> One can meet all ***possible*** relevant non-metaphysical objections to *p*.
> (1973, 75)

Ackermann observes that the first of these readings is incompatible with the Ideal Analysis, as one might correctly assert knowledge despite there existing some future objection that ultimately proves p to be false (ibid, 76). However, at this point Ackermann does something quite unexpected. Rather that simply discarding the first read on the grounds that it conflicts with the truth condition on knowledge, he instead proposes the Pragmatic Analysis, which is compatible with this read. He summarizes this analysis in the following way:

> There are several points to notice about the Pragmatic Analysis. Suppose *Kap*[5] can be asserted. It is compatible with this that ~*p* is true. Now suppose that some new relevant non-metaphysical objection to *p* is discovered which can't be met. (In a sense, this raises the possibility that, compatibly with

[4] It's unclear here whether by using the word "much" Ackermann means to indicate that this argument has received some, but little, attention, or if he's unaware of any such attention but wants to play it safe. At no point in the monograph does Ackermann reference any previous attempt to argue against factivity.

[5] a knows that p

> current knowledge, ~*p* might be true.) *Kap* is no longer true. Should ~*p* now be discovered true, then *Kap* is certainly not true. We have the same situation that we had on the Ideal Analysis, ~*p* and *Kap* cannot be jointly asserted by a rational agent. The different in the situation is provided by involvement with temporality. On both analyses, once ~*p* is found out to be true, one must conclude that a rational agent can't know that *p*. The difference comes in connection with our attitude towards past knowledge claims. On the Ideal Analysis, if ~*p* is true, one could never have known that *p*. On the Pragmatic Analysis, the situation is somewhat different. (ibid, 76)

Here we might make a few key observations about Ackermann's argument. First, he clearly assumes that the assertability of "S knows that p" to entail that S knows that p. This is a perfectly reasonable assumption, as it is widely accepted that assertability entails truth (see Turri 2016). This allows Ackermann to use our judgements that knowledge assertions in certain cases are acceptable as evidence for knowledge in those cases, which is the primary method he employs. Next, and most significantly, he takes there to be legitimate cases of knowledge assertion even when it turns out the known proposition wasn't actually true. The following is one example he provides:

> I used to know that if I pressed button A, the temperature would reach safe levels within ten minutes. (1973, 77)

Ackermann imagines this might be asserted by a scientist in response to the "unexpected performance" of an apparatus (ibid, 77). Even in the case that the scientist's belief turned out to be false, he considers the above perfectly assertable, which entails that the scientist really did have knowledge. Thus, the Pragmatic Analysis seems to capture "facts about the way in which we ordinarily conceive of knowledge" that factive analyses cannot (ibid, 77). Ackermann goes on to argue that Pragmatic Analysis is more explanatory than factive analyses, especially in ordinary cases in which we assert knowledge despite not having access to all relevant alternatives (ibid, 108-115). In short, Ackermann takes his non-factive Pragmatic Analysis of knowledge to be explanatorily successful in a broader range of cases than the factive Ideal Analysis.

Although noteworthy, Ackermann's argument is obviously flawed. The most prominent problem is that he confuses the acceptability of utterances for their literal assertability. Yes, we will often say things as in the "temperature button" example, stating that we used to know a proposition that later proved to be false, and there seems to be nothing unacceptable about this practice. However, it is a mistake to think that we literally assert knowledge in these cases. Rather, we are instead asserting along the lines that we believed that p or thought we knew that p, not that we actually ever had knowledge that p. This becomes especially obvious if we move away from our linguistic judgements regarding utterances, and instead simply look at our *epistemic* judgements. If we consider whether the scientist ever actually knew that if he pushed button A then the temperature would reach safe levels within ten minutes, given that that wasn't true, our intuitive judgement is clearly that he did not. Moreover, while Ackermann is certainly able to provide examples in which "knows" is used in a non-factive way, we don't actually judge that S has knowledge in these cases.

The obvious lesson here is that while sometimes "knows" is used in a non-factive way, this does not itself motivate a departure from the factive analysis of knowledge. For this, we need cases in which we actually judge a false belief to be knowledge.[6] We'll see something quite similar with the "protagonist projection" discussed in the following section.

2.3 Hazlett's "The Myth of Factive Verbs"

Easily the most familiar contemporary objection to the factivity of knowledge is Allan Hazlett's 2010 "The Myth of Factive Verbs." As with Ackermann, Hazlett's argument is based on non-factive uses of "knows" in ordinary language. However, Hazlett's account differs from Ackermann's in two major respects. First, Hazlett does not maintain that the epistemological concept of knowledge is non-factive. Unlike Ackermann, he doesn't propose, or even suggest, a non-factive analysis of knowledge. Second, upon its publication Hazlett's argument attracted a significant amount of attention. Unlike Ackermann, Hazlett's account is well known, spawning a number of objections in the literature. My review of Hazlett will follow these two points. First I'll introduce his argument, after which I'll discuss one of the major objections that has been raised against this argument. In doing so, I'll contrast Hazlett's claims with the ones I'll advance here. Distinguishing between Hazlett's claims and my own serves a valuable clarificatory role for this project. As both our arguments and conclusions are almost entirely distinct, it would be a mistake to consider this project to be a continuation of, or spiritual successor to, the "Myth of Factive Verbs" argument.

Hazlett's account is based on seemingly acceptable utterances that use purportedly factive verbs in non-factive ways. These verbs include "learns," "remembers," and "realizes," but here I'll discuss his argument only as it pertains to "knows." Hazlett maintains utterances like the following seem acceptable, despite using "knows" (or "knew") in a non-factive way:

> Everyone knew that stress caused ulcers, before two Australian doctors in the early 80s proved that ulcers are actually caused by bacterial infection.
>
> He figures anything big enough to sink the ship they're going to see in time to turn. But the ship's too big, with too small a rudder … it can't corner worth [expletive]. Everything he knows is wrong. (2010, 501)[7]

On the basis of the seemingly acceptable non-factive usage of "knows" and "knew" in ordinary language, Hazlett argues that the ordinary concept of knowledge is non-factive (ibid, §4). Here I'd like to note that this is one major difference between my argument and Hazlett's, as I make no use of intuitive

[6] Of course, one might also make the argument on other theoretical grounds, such as avoiding a skeptical risk. As already mentioned, I employ this sort of argument alongside one based on actual judgements of knowledge for false belief.

[7] As cited by Hazlett, the first example is adapted from Achenbach (2005), and the second is Brock to Bodine in *Titanic (1997).*

judgements about the acceptability of utterances, or any other sort of linguistic data. However, the most significant difference between Hazlett's account and my own is that Hazlett doesn't think that the epistemically interesting concept of knowledge is non-factive. In fact, he maintains that it is indeed factive, and argues on this basis that the ordinary concept of knowledge that underlies "knows" is distinct from epistemological concept of knowledge. He summarizes this argument in the following way:

> If I'm right, then epistemologists may have reason to stop looking at linguistic phenomena altogether—at least if they want to keep working on anything like the standard analysis of knowledge. The concept of knowledge that epistemologists have been interested in since the *Meno* is a factive concept (in the sense that nothing false can be known). But, if I'm right, the concept of knowledge that serves as the meaning of 'knows' in ordinary talk isn't. This is strong prima facie evidence that traditional epistemology shouldn't be especially interested in the concept of knowledge that serves as the meaning of 'knows' in ordinary talk. . . For the *episte*mologist is interested in an *epistemic* concept of knowledge, if she is interested in a concept of knowledge at all. What I'm claiming is that epistemologists have every right to insist that knowledge (as they understand it) is factive—but the price to pay for this (which many will be happy to incur) is to give up the linguistic method described above. I'm suggesting, in other words, a divorce for the linguistic theory of knowledge attributions and traditional epistemology. (ibid, 499-500)

As the aim of this thesis is to argue that the epistemologically interesting concept of knowledge is non-factive, and on this basis advance the non-factive analysis *of this concept*, we might observe that my project diverges from Hazlett's at a fundamental level. However, this doesn't mean that there is nothing to be learned from Hazlett' argument for the non-factive use of "knows." Here there are two aspects of his account I want to discuss. I want to focus primarily on one key objection that has been levied against his account, as it highlights a certain pitfall that an argument against factivity might fall into. Before this, however, I first want to discuss Hazlett's take on the contradictoriness argument mentioned in 2.1.

First, Hazlett focuses specifically on the first-person utterance "I know p, but not p," maintaining that it is not contradictory, but Moore-paradoxical[8]:

> 'I know p, but not-p' is not contradictory, but an utterance of it is Moore paradoxical – to know that p is to believe that p, and 'I believe p, but not-p' is paradigmatically Moore paradoxical. It is possible to mistakenly take a sentence, the utterance of which would be Moore paradoxical, for a contradiction. An utterance of 'I know p, but not-p' is always improper, but the sentence is not a contradiction. (ibid, 506)

Beyond this, Hazlett also offers a more general account, on which "S knows that p" often implies the truth of p, but doesn't entail it generally. Following Grice's (1989)

[8] The operative idea here is simply that it seems paradoxical to say things of the form "I believe that p, but not p," but it isn't contradictory. For a detailed account of Moore's paradox see Wittgenstein (1953, IIx).

account of implicature, he begins by assuming conversational cooperation for our utterances involving "knows:"

> The explanations I prefer are all more or less Gricean in flavor. I presuppose the following: that it is mutually assumed by speakers the people generally conform to a principle of conversational cooperation, which requires general, and mutually assumed, conformity to at least three maxims: that of Quality ('Do not say anything you believe to be false, or which you don't have reason to believe is true'), that of Quantity ('Make your contribution to a conversation as informative, and only as informative, as is required'), and that of Relation ('Make your contribution to a conversation relevant').
> (ibid, 511-2)

Beyond this, he also assumes that attributing knowledge to S entails attributing epistemic warrant to S. Taken together, these two assumptions allows us to explain how knowledge attributions might imply truth. For example, if I say "S knows that she's drinking coffee," this entails an attribution of sufficient epistemic warrant to S for believing that she is drinking coffee. This means that if I have any reason to suspect that S really isn't drinking coffee, it would be a violation of the maxim of quantity to withhold it. Therefore, by attributing knowledge to S without mentioning any contrary evidence, I imply that S's belief is true. In this way, Hazlett provides a way in which we might explain the seeming contradictoriness of utterances like "S knows that p, and not p" without positing factivity. While I will argue later that there is an even simpler explanation for this seeming contradiction (chapter 4, section 4), here it is worth noting how Hazlett's account pushes back on the idea that factivity plays an indispensible explanatory role for these cases.

To close my discussion of Hazlett I'd like to discuss one key objection that has been raised against his account.[9] In doing so, I hope to highlight a major pitfall associated with the denial of factivity, which my argument will take care to avoid.

One of the requirements of Hazlett's argument is that non-factive uses of "knows" in ordinary language are literal. If the speakers of utterances, as in the examples above, don't literally mean "knows" or "knew" to be non-factive, we of course don't need to posit a non-factive concept of knowledge to explain these usages. However, one major challenge for Hazlett's account is that it does not appear that "knows" is actually used literally in the types of cases he cites. As argued by Buckwalter, they can instead be explained as cases of "protagonist projection" (2014a, 391). First described by Holton (1997), protagonist projection refers to the practice of using sentences to non-literally convey the beliefs of others:

> I suggest that these sentences work by projecting us into the point of view of the protagonist; let us call the phenomenon *protagonist projection*. In each case the point of view into which we are projected involves a false belief. We describe the false belief using words that the protagonists might use themselves, words that embody their mistake. So we deliberately use words in ways that do not fit the case. (ibid, 626)

[9] Note that I won't discuss objections that are highly specific to Hazlett's argument, such as those from Turri (2011) and Tsohatzidis (2012). See Hazlett (2012) for his replies to these objections.

Protagonist projection offers a compelling explanation of the non-factive usage of "knows." When we say that what someone "knows" is wrong, we are projecting ourselves into the perspective of the protagonist, using words she might use, not literally maintaining the protagonist actually knows. Beyond simply offering this as a rival explanation for Hazlett's observations, Buckwalter conducted a series of empirical studies probing the intended meaning of the utterances that provide the basis for Hazlett's argument. For example, in one experiment Buckwalter asked participants whether Hazlett's "ulcer" example meant

> (A) Everyone thought they knew.
>
> (B) Everyone really did know. (2014a, 397)

Participants overwhelmingly selected option (A), indicating that they did not take this utterance to literally mean that people once had knowledge of a falsehood (ibid, 397). Rather, these non-factive uses of "knew" and "knows" appear to be best understood as cases of protagonist projection. The speakers in these cases are simply using non-literal language to convey the mistaken beliefs of the protagonists, specifically that they have knowledge.

Given the role protagonist projection might play in the acceptability of the non-factive usage of knows, it is important that my account avoids this problem. The primary way I do this is by avoiding linguistic data entirely. As my argument is not based on the usage of "knows" in language, it is difficult to object to it via protagonist projection. It is conceivable that one might attempt to adapt this objection to the apparent *epistemic* judgements that false beliefs are known proposed in the course of my argument. However, for this reason, among others, my argument is not based solely on our intuitive judgements. As mentioned before, it is also comprises of a key anti-skeptical component. In this manner, it is able to circumvent the objection that an unidentified variant of protagonist projection plays a role in my argument.

2.4 Pritchard and Kelp: "Weakening" Factivity

As a potential response to the "paradox of knowability," Fitch's (1963) argument that all truths being knowable (along with a few plausible assumptions) paradoxically entails that all truths are known, Pritchard and Kelp briefly entertain the option of "weakening the factivity principle" (2010, 329). Although they quickly conclude that this route is unsuccessful, I mention this here because I also commit to something in my argument that I call "weakening factivity," and I want to clarify that I don't at all mean by this what Pritchard and Kelp do. I use "weakening factivity" interchangeably with "replacing factivity with a weaker constraint on knowledge." Crucially, as this is done with the explicit purpose of counting certain false beliefs as knowledge, this means that falsehoods can be known. Conversely, on their version of weakening, while knowledge does not entail true, falsehoods still cannot be known:

> There is a potential logical gap—at least by semantic anti-realist lights—between the claim that there cannot exist any cases in which an agent knows a falsehood and the factivity claim that knowledge entails the truth of the proposition known. (ibid, 330)

Again, here I'll argue that there are certain, in fact many, false beliefs that we should still count as known. Accordingly, my "weakening" of factivity goes much further than Pritchard and Kelp have in mind here.

3. Method: Analytic and Naturalistic Components

In this section, I'll discuss the philosophical method employed by this project, and in doing so clarify a few of the basic assumptions it will commit to. At least cosmetically, the method I'll employ might be divided into two components: (1) the analytic component and (2) the naturalistic component. The analytic component is quite standard, involving the testing of proposed necessary conditions on knowledge against our intuitive judgements and basic theoretical requirements. The naturalistic component is more novel, most notably in its conclusion that even the most fundamental questions of analytic epistemology in fact reduce to empirical questions. Accordingly, while I'll discuss both the analytic and naturalistic components in this section, I'll devote significantly more time to the latter.

3.1 The Analytic Component

This project is analytic in the sense that it is fundamentally interested in the analysis of the concept of knowledge, primarily in whether factivity is a necessary condition on knowledge. While it so happens that this analysis will involve a number of empirical findings, it is important to note that I take knowledge to be a *concept*, not a natural kind (see §3.2). Here I want to briefly highlight the analytic nature of the two primary arguments advanced in this project. These are (1) the argument in chapter three in favour of rejecting the factivity condition on knowledge and (2) the argument in chapter four that the non-factive analysis of knowledge is more successful that the factive analysis. I'll discuss the analytic features of each argument in turn.

In chapter three, I argue that we should drop the factivity constraint on knowledge in favour something weaker. Both the nature of this claim and the argument I present for it are unmistakably rooted in the analytic project. First, the claim itself is about how best to analyse the concept of knowledge. On this claim, knowledge is best analysed as a non-factive concept, and epistemologists have made a mistake in assuming that knowledge is a factive concept. Moreover, the argument I provide for this claim is equally analytic in nature. This argument is largely conceptual, and rests on familiar assumptions of analytic epistemology:

1. Necessary conditions on knowledge should track our intuitive judgements about whether an agent has knowledge.

2. Necessary conditions on knowledge shouldn't invite skepticism.

This first assumption comes with the disclaimer that at least sometimes we observe troublesome patterns of intuitive judgement that we may not want necessary conditions on knowledge to track. I discuss this in chapter three (§5). Nevertheless, in my argument against factivity, I maintain that not only does it seem that we judge certain false beliefs to be known, but that these judgements are certainly not like those we might exclude from the analysis of knowledge. This second assumption is fundamental to this project: Necessary conditions on knowledge should not entail that large swaths of putative knowledge are in fact unknown. As I will discuss throughout the following chapters, if one is comfortable with skepticism, many of my conclusions might avoided.

Chapter four builds on the analytic endeavour started in chapter three by outlining what the non-factive analysis of knowledge might look like. Again, it assumes that the purpose of such an analysis is to track our intuitive judgements without giving way to skepticism. Beyond this, I don't think there's anything else interesting to say at this point.

3.2 The Naturalistic Component

As "naturalism" in epistemology is used to describe a general inclination towards reducing the distance between epistemological theory and empirical science (see Rysiew 2017 for a review), my project can certainly be classified as committing to a sort of naturalism. Here I hope to make clear the exact commitment to naturalism I have in mind by contrasting the naturalism of this project with a few foundational naturalistic accounts of epistemology. Before doing this, however, I first would like to provide a preliminary overview of the basics of the naturalistic component of my account. Most importantly, it is important to note that I do not take naturalism to be something that is in competition with conceptual analysis, or even something that is extra-analytic. Rather, I think we might best understand empirical findings as a tool indispensible in the practice of analysis. Contemporary psycho-cognitive research is highly relevant to fundamental analytic questions, like whether knowledge is factive, and thus these questions cannot be properly adjudicated without taking this research into consideration. If other conceptual assumptions are largely agreed upon, this can mean that fundamental analytic questions can reduce to empirical ones. This is perhaps the most important lesson to be gained from this project. Even more significant than the singular claim that knowledge is non-factive, this project suggests that even the most fundamental philosophical claims are subject to empirical critique and revision.

One way we might cash out this lesson is that empirical research can establish key premises in conceptual arguments. This of course is not at all unprecedented. An excellent illustration of how empirical findings can ground conceptual arguments comes from the debate whether knowledge-that and knowledge-how are conceptually distinct. One of the major arguments that knowledge-how does not reduce to knowledge-that uses empirical observation to support the premise that we might retain knowledge-how even after losing any associated knowledge-that (Wallis 2008; Adams 2009; Brown 2013). My intention is to do something similar here, but on a larger scale.

Beyond involving empirical research in some substantive way, it is not widely agreed upon what makes an epistemological account "naturalistic." Here I will discuss three of the most familiar articulations of naturalism in epistemology, those of Quine, Kornblith, and Goldman. Beyond simply providing an overview, I'll also contrast each which the version of naturalism I'll advance.

3.2.1 Quine's "Epistemology Naturalized"

Modern discussion of naturalism in epistemology begins with Quine's 1969 essay "Epistemology Naturalized." Here I'll primarily discuss Quine's core claims and contrast his version of naturalism with my own, along with the motivations of and objections to his account.

The defining feature of Quine's naturalism is its radical reimagining of the relationship between epistemology and natural science. Going well beyond the idea that epistemological theorizing might benefit from input from the empirical sciences, Quine recommends that epistemology be subsumed by the natural sciences entirely, assuming a position as a sub-field of psychology:

> Epistemology, or something like it, simply falls into place as a chapter of psychology and hence of natural science. It studies a natural phenomenon, viz., a physical human subject. This human subject is accorded a certain experimentally controlled input—certain patters of irradiation in the assorted frequencies, for instance—and in the fullness of time the subject delivers as output a description of the three-dimensional external world and its history. The relation between the meagre input and the torrential output is a relation that we are prompted to study for somewhat the same reasons that always prompted epistemology; namely, in order to see how evidence relates to theory, and in what ways one's theory of nature transcends any available evidence. (1969, 82-3)

Quine's motivation for this proposal lies largely in his view that "epistemology is concerned with the foundations of science" (ibid, 69). The idea here is that if we want to give an account of how sense data can ground scientific theories, the best way to do this is to empirically observe how sense data *actually* results in scientific theory. I think there's a lot to like about this idea, which I'll discuss shortly. Before this, however, we should address the shortcomings of Quine's account.[10] First, there is the glaring issue that Quine's conception of epistemology was, even at the time of his essay, obsolete. The epistemological program that Quine uses as a framing device is "rational reconstruction," Carnap's (1928) endeavour to "account for the external world as a logical construction of sense data" (1969, 74), which as noted by Kelly, "had already been abandoned" when Quine wrote *Naturalized* (2014, 24). Accordingly, whatever the extent to which my naturalism might agree with Quine's, it is obvious that I will not share in his motivations.

Next, perhaps the greatest shortcoming of Quine's account is that it eliminates epistemology as a philosophical discipline entirely, recasting it as a

[10] For a review of the various objections raised against Quine's account, as well as responses, see Rysiew (2017).

psychological one. This famously raises a serious problem with epistemic normativity, as it seems that psychology, and thus psychologically constrained epistemology, is wholly descriptive (Kim 1988; for a response see Quine 1992, 19). However, what concerns me more about the idea of reducing the discipline of epistemology to psychology is that epistemological concepts and psychological concepts often vary considerably. Perhaps the best example of this is the concept of knowledge itself, which as a technical psycho-cognitive concept is largely unrecognizable to epistemologists. This is simply because epistemology and psychology want a concept of knowledge to do different things. Roughly, psychology, as well as other related scientific disiplines, needs a concept that can capture a broad range of experimental data, whereas epistemology needs a concept that can capture a broad range of epistemic judgements (see chapter four, section three for more on this). Were psychology to annex epistemology, presumably this would mean abandoning all those aspects of epistemology that are not psychologically interesting. My suspicion is that the analysis of knowledge would be the first thing to go. As I have no interest in this, my project has no intention of following the Quinean reduction of epistemology to psychology.

A final contrast I want to draw between my account and Quine's is that I don't take epistemology to study a "human subject," or any "natural phenomenon." As will come up again in the next sub-subsection, I do not take knowledge to be a natural kind. This means that, although empirical research can be extremely useful, the study of knowledge cannot be a fundamentally empirical one.

Despite the many ways in which my project is not at all "naturalized" in the Quinean sense, there is one key feature they share: They both recognize that epistemological questions might reduce to empirical ones. For Quine, this question was how sense data provides a foundation for science. For my project, this question is whether knowledge is factive. Whether it is a factive or non-factive concept that best fulfils the basic theoretical requirements of the analysis of knowledge (capturing epistemic judgements and avoiding skepticism) is contingent about specific facts about how our cognitive systems represent information.

To close, I want to note an especially interesting point of contact between Quine and this project, as Quine seems to foreshadow what I discuss as claim (2) in chapter two:

> Our retinas are irradiated in two dimensions, yet we see things as three-dimensional without conscious inference. Which is to count as the observation—the unconscious two-dimensional reception or the conscious three-dimensional apprehension? In the old epistemological context the conscious form had priority, for we were out to justify our knowledge of the external world by rational reconstruction, and that demands awareness. Awareness ceased to be demanded when we gave up trying to justify our knowledge of the external world by rational reconstruction. What to count as observation now can be settled in terms of the stimulation of sensory receptors, let conscious fall where it may. (Quine 1969, 84)

As we'll see in chapter two, this is a very rudimentary account of the multiple references frames in which our cognitive systems represent spatial information.

Nevertheless, Quine manages to not only highlight the basic idea that mental representation is not unitary, but he also anticipates the challenge my account will face in basing an epistemological account on content that doesn't rise to full conscious awareness (see chapter 3, section 3). At least in this sense, we might say that I, with Quine, will reject the "old epistemological context" on which "conscious form had priority." All the false beliefs I argue are known have, to a significant extent, content that falls below the level conscious awareness.

3.2.2 Kornblith: Knowledge as a Natural Kind

Next, I'll discuss the version of naturalism espoused by Kornblith, after which I'll contrast it with my own. As with Quine, while there are certain areas of overlap, the methodology I employ in this project differs substantively from that recommended by Kornblith.

Unlike Quine, Kornblith's framing of epistemology corresponds with the contemporary significance placed on the study of knowledge. However, he still maintains a quite heterodox stance on what this object of study actually is. Kornblith rejects the idea that epistemology should be interested in the concept of knowledge at all, thereby writing off the analytic project entirely:

> My insistence that epistemology should not concern itself with our concept of knowledge requires that I depart, in important ways, from some common practices. I will not, for the most part, be comparing my account of knowledge with my intuitions about various imaginary cases; I will not be considering whether we would be inclined to say that someone does or does not have knowledge in various circumstances. I do not believe that our intuitions, or our inclinations to say various things, should carry a great deal of weight in philosophical matters. (2002, 1-2)

In place of the study of a concept, Kornblith suggests that epistemology transform itself into the study of the *phenomenon* of knowledge, thereby casting knowledge as a natural kind (emphasis mine):

> Where should we turn, and how should we proceed, if we are to investigate the phenomenon of knowledge itself? It is not necessary to found a new discipline here, for there is already a good deal of empirical work on the subject. One of the more fruitful areas of such research is cognitive ethology. There is a large literature on animal cognition, and workers in this field typically speak of animals knowing a great many things. They see animal knowledge as a legitimate object of study, a phenomenon with a good deal of theoretical integrity to it. Knowledge, as it is portrayed in this literature, does causal and explanatory work . . . If cognitive ethologists are even roughly right, then talk of animal knowledge is not a mere *façon de parler*; rather, there really is such a thing as animal knowledge. ***Knowledge constitutes a legitimate scientific category. In a word, it is a natural kind.*** (ibid, 28-9)

At this point there are several key ways in which we might contrast the naturalism of this project with that presented by Kornblith. First, I of course don't see naturalism as in any way supplanting the analysis of knowledge, but rather playing a crucial role within the analytic project. A second, very much related contrast is that I don't take knowledge to be a natural kind, being committed to the study of knowledge *the concept*.

Beyond these coarse differences of method, there is a more layered and interesting contrast to be made. Like Kornblith, I think the scientific use of "knowledge" is epistemologically interesting and ought to be taken seriously. Indeed, in chapter four I formulate an analysis of knowledge that seeks to unify the epistemological and scientific ("psycho-cognitive") concepts of knowledge. In this way, we might explain how the concept of knowledge might play dual epistemological and scientific roles. Here we might again note a significant contrast with Kornblith. First, I am ultimately still interested in concepts of knowledge, while Kornblith is not. Second, the empirical literature Kornblith focuses on, cognitive ethology, is rather unlike anything you'll see in this project, which for the most part fits under the broad purview of "cognitive neuroscience" (discussed in the next section).

There is a final contrast I'd like to discuss here, which relates to the final account of knowledge that Kornblith endorses. Despite not being interested in the concept of knowledge at all, Kornblith comes to the familiar conclusion that knowledge is "reliably produced true belief" (ibid, 28). Remarkably, it appears that the phenomenon of knowledge might be described via reliabilism, a well-trodden theory of the concept of knowledge. Not being an expert in cognitive ethology, I cannot speculate as to whether this accurately captures how "knowledge" is used within that field. However, I can note such an account doesn't even come close to capturing the explanatory roles the technical concept of knowledge plays in empirical literature generally. Again, I discuss this in chapter four, section three, focusing primarily on the usage within cognitive neuropsychology. While we might concur with Kornblith that the psycho-cognitive concept of knowledge is robust, doing real theoretical work, it clearly doesn't come with a belief or truth condition. Traditional epistemological theory must make significant concessions in order to capture the myriad of different cases in which psycho-cognitive knowledge plays explanatory and predictive roles. "Reliably produced true belief" simply doesn't cut it.

In contrasting this project with Kornblith's naturalism, I hope to have highlighted a few key features of my take on naturalism. Most interesting here is the idea that the technical concept of knowledge we find in empirical literature is not identical with the standard epistemological concept, nor might it be captured with a standard epistemological account. However, this doesn't mean that this concept is not of epistemological interest.

3.2.3 Goldman's "Scientific Epistemology"

Far more so than either the respective versions presented by Quine and Kornblith, Goldman's naturalism, which he dubs "scientific epistemology" (1993a), somewhat closely matches the naturalism of my project. Here I'll highlight some of the central claims of Goldman's account and again contrast them with my own.

To begin with, I should note that Goldman's naturalism extends far beyond epistemology (see 1992; 1993b). However, here I'll stick to his epistemological naturalism. Put simply, the "scientific epistemology" proposed by Goldman is far more conservative than the previously discussed proposals from Kornblith or (especially) Quine. To begin with, not only does Goldman frame epistemology in terms of the study of knowledge (contra Quine), but also he takes conceptual analysis to be the appropriate means of studying knowledge (contra Kornblith[11]). Goldman is notable in this respect as the only naturalist I'll discuss here to have (famously) advanced multiple analyses of knowledge over the course of his career (1967; 1993a). He views the analysis of knowledge as a largely *a priori* conceptual project, which might benefit from modest empirical input, articulated clearly in the opening to his 1988 paper "Psychology and Philosophical Analysis:"

> It is often said that philosophical analysis is an *a priori* enterprise. Since it prominently features thought experiments designed to elicit the meaning, or semantic properties, of words in one's own language, it seems to be a purely reflective inquiry, requiring no observational or empirical component. I too have sometimes acquiesced in this sort of view. While arguing that certain phases of epistemology require input from psychology and other cognitive sciences, I have granted that the more 'conceptual' stages of epistemology are strictly philosophical and (hence) non-empirical. In this paper I want to qualify this position. I shall suggest that psychological theories can have a bearing on philosophical analysis; they can support the plausibility or implausibility of specific analyses. (195)

Goldman imagines a very narrow role for psychology here, as by "support the plausibility or implausibility of specific analyses" all he means is that psychology can provide a broader and more robust account of our intuitive judgements than can individual introspection. He specifically argues for a relevant alternatives account of knowledge based on the fact that this is most congruous with "cognizers' judgements and categorizations" (ibid, 199), as evidenced by psychological observation. The idea here is that since psychology might describe factors that contribute to our judgements generally, it can specifically help describe factors that contribute to our judgements about knowledge.

It is here I'd like to draw a first contrast between Goldman's naturalism and my own. While I too am interested in the analysis of knowledge, I think empirical evidence might do more analytic work than Goldman allows for. First, I think this evidence, emerging more from cognitive neuroscience than psychology proper, might highlight not merely previously hidden facts about our judgements, but rather previously hidden facts about cognition about which we might form new judgements. For example, as I argue in chapter three, empirical evidence highlights previously unconsidered cases, about which it appears we judge false beliefs to be known. Now of course it is conceivable that these cases might have been concocted as putative counterexamples in the absence of any empirical observation. However, it is not only unlikely that they might have been rendered as richly as via cognitive neuroscience, but I think the very fact these cases are actual,

[11] He explicitly dismisses the idea of knowledge as a natural kind as "dubious" (1988, 195).

and in fact ordinary, sways our judgements. Beyond this, I think empirical evidence does a superb job of highlight skeptical problems, as it lends itself to the actual description of our actual knowledge. Thus, if, say, neuroscientific evidence suggests that most all of our ordinary knowledge has feature X, we can be sure to avoid any analysis of knowledge on which feature X is taken to preclude knowledge.

Goldman's naturalism is of course not limited to the analysis of knowledge, and he takes the primary epistemological relevance of empirical observation to fall outside the analytic project. Instead, Goldman sees cognitive science as especially valuable in answering the question, "What are the sources and prospects for human knowledge and rationality?" (1993b, 1). This wouldn't be all that relevant for our purposes, except that Goldman maintains that empirical observation can "help *set* the standards for rationality" (ibid, 31). Here, the idea is that evidence from cognitive science can help to not only explicate what it means for an agent to possess evidence, but also give us a "realistic perspective" on the actual extent to which we ordinarily possess evidence (ibid, 31).

The reason I draw attention to Goldman's claim that empirical evidence can help set the standards for rationality is that I'll argue along similar lines for standards of knowledge. Specifically, in chapter four, I'll argue in favour of replacing the factivity condition with a "truthlikeness" condition on knowledge, on which beliefs can be known only if there are very close to being true. As, "very close to being true" is incredibly vague, I offer that the level of truthlikeness necessary for knowledge might be set by the level of accuracy of the cognitive processes that produce knowledge. In this manner we might understand how the standards of knowledge are, at least partially, an empirical matter, as we need to observe the specific cognitive processes that produce knowledge in order to determine the level of truthlikeness necessary for knowledge.

A final similarity between my naturalism and Goldman's can be found in the epistemological significance he places of folk practice:

> Whatever else epistemology might proceed to do, it should at least have its roots in the concepts and practices of the folk. If these roots are utterly rejected and abandoned, by what rights would the new discipline call its 'epistemology' at all? (1993a, 271)

Perhaps even more so than Goldman, who takes these "folk pathways," as he calls them, to be subject to revision (ibid), this project places a substantial value on our epistemic folk practices. Specifically, I put preserving our ordinary practices of belief and knowledge -attribution as a key epistemological priority, and argue for the preservation of these practices even at the cost of factivity.

In short, while Goldman offers an account of naturalism somewhat milder than mine, his is easily the closest of the three discussed here. Most noteworthy is his commitment to the analytic project and his suggestion that empirical evidence has a role to play in analysis. The arguments I'll present in the following chapters not only highlight the empirical component of the analysis of knowledge, but also suggest that, if anything, Goldman was too conservative in his assessment of how relevant empirical observation can be.

4. Cognitive Neuroscience

In this section, I want to briefly attempt to introduce and categorize the types of empirical evidence I'll utilize throughout this project (primarily in chapters two and five). As I provide discussion of specific details as they become relevant, here I'll limit discuss to those broad generalities necessary to understand the specific details. The best way I can think to do this is by simply providing a brief overview of cognitive neuroscience as a discipline. This isn't to say that all the empirical research I'll cite falls nicely under the label "cognitive neuroscience, rather only that the (i) methods of cognitive neuroscience exhaust the empirical methods I'll discuss, (ii) this project will view observations produced by these methods through the lens of cognitive neuroscience, and (iii) the primary theoretical marker of cognitive neuroscience, the "multiple levels" approach to inquiry, closely correlates with the multi-level arguments I advance in subsequent chapters. Accordingly, we might understand that my goal with this section is to provide a rough overview of the empirical contents of this project, not describe a topography on which we might easily map each of the hundreds of individual studies we'll encounter.

Cognitive neuroscience is, as put by Boone and Piccinini, the "new mainstream approach to studying cognition" (2016, 1509). Here I'll provide a brief overview of its approach and empirical methods, especially as it relates to my project. As a discipline, cognitive neuroscience is defined largely by a "multiple levels of analysis" approach to the study of human cognition (Ochsner & Kosslyn 2013, 2). This approach has been categorized as supplanting that of traditional cognitive science, with its conspicuous lack of interest with descriptions at the neural/mechanistic level (Boone and Piccinini 2016). The multi-level approach of cognitive neuroscience was pioneered by vision scientist David Marr, who argued:

> To understand what vision is and how it works, an understanding at only one level is insufficient. It is not enough to be able to describe the responses of single cells, nor is it enough to be able to predict locally the results of psychophysical experiments. Nor it is enough even to be able to write computer programs that perform approximately in the desired way. One has to do all these things at once. (1982, 329)

Here I don't want to suggest that either my project or cognitive neuroscience as a whole is committed to the specifics of Marr's computational approach. The point here is simply that this commitment to multiple levels of analysis is the definitive feature of cognitive neuroscience. The result is theory that encompasses description at the behavioural, psychological, and neural levels (Ochsner & Kosslyn 2013), which affords it certain advantages over parallel approaches:

> The resulting multilevel descriptions have many advantages over the one- or two-level accounts that are typical of traditional approaches in allied disciplines such as cognitive psychology. These advantages include the ability to use both behavioral and brain data in combination—rather than just one or the other taken alone—to draw inferences about psychological processes. In so doing, one constructs theories that are constrained by, must

> connect to, and must make sense in the context of more types of data than theories that are couched solely at the behavioral or at the behavioral and psychological levels. (ibid, 2)

The multi-level approach that defines cognitive neuroscience, or at least something resembling it, is also a key feature of this project. Epistemological analysis is of course the "top" level, which is informed by observation at the behavioural level, which is in turn informed by more direct observation of the brain itself. We see this illustrated most dramatically in section three of chapter five, in which an epistemological argument against the factivity of knowledge is grounded in empirical descriptions of the representational content of individual neurons, with observations at the behavioural level also playing a key intermediary role. In a sense, one might think of my approach as "tacking on" the epistemological level of inquiry to the extant cognitive-neuroscientific levels, thereby grounding epistemology in the multi-level empirical descriptions provided by cognitive neuroscience. In short, this is the idea at the heart of my project: Not only can standard epistemology accommodate a sort of multi-level naturalism, but epistemological analysis stands to benefit greatly from it.

As cognitive neuroscience adheres to a multi-level approach to cognition, it incorporates multiple types of empirical evidence, including behavioural, neuroimaging, neuropsychological, and neurophysiological (Ochsner & Kosslyn 2013). As the first two of these are likely familiar already, they require little explanation. Behavioural studies typically look at the behaviour of healthy individuals to draw conclusions about our cognitive architecture. I discuss specific techniques used in behavioural studies as they become relevant (see chapter two). It's likely you're also already oriented with the concept of neuroimaging studies, which "use functional imaging methods to describe the network of processes active in the healthy brain when engaged in a particular behavior" (ibid, 3). As the details of functional neuroimaging techniques aren't relevant for our purposes, I won't discuss them here.

Neuropsychological evidence is that associated with the techniques of cognitive neuropsychology, a predecessor and rival to cognitive neuroscience that now has largely been subsumed by it (Ochsner & Kosslyn 2013). As with cognitive neuroscience, cognitive neuropsychology utilizes data from neuroscience to inform theories of cognition (Caramazza 1992). However, unlike cognitive neuroscience, the only neuroscientific data of interest to cognitive neuropsychology comes from "patterns of performance produced by brain-damaged subjects" (ibid, 1980). Again, I'll discuss the specifics as they become relevant, but the core idea here is that brain lesions can impair specific cognitive systems while leaving others intact:

> By studying the ways in which behavior changes as a result of the unhappy accidents of nature (e.g., strokes, traumatic brain injuries) that caused lesions of language areas, memory areas, and so on, we can discover the processing modules that constitute the mind. (Ochsner & Kosslyn 2013, 2)

Any evidence that indicates specific cognitive architecture on the basis of these lesion studies is what we might call "neuropsychological evidence."

Finally, there is neurophysiological (or "electrophysiological") evidence. As with neuroimaging, neurophysiological studies observe neural activity given specific inputs to draw conclusions about cognitive architecture. However, rather than relying on functional imaging, neurophysiological studies use electrodes implanted in the brain to directly measure the activation of either single neurons or neural populations (Carpenter 1990, 14). This is especially useful in identifying the neurons and neural populations that represent specific stimuli, along with the nature of these representations (ibid).

The empirical evidence I present in the argument against factive knowledge spans behaviour, neuroimaging, neuropsychology, and neurophysiology. In this way too we might understand the naturalistic method of this project to mirror that cognitive neuroscience. Again, to be clear, this doesn't mean that all the empirical research I cite was explicitly conducted under the banner of "cognitive neuroscience."

Chapter Two: Empirical Foundations

As discussed in the previous chapter, my argument against the factivity of knowledge rests in large part on empirical claims. In this chapter, I discuss the following four claims regarding mental content, which form the foundation for my account:

(1) *Perception can occur in the absence of consciousness.*

(2) *The same information is often represented internally in a multitude of ways.*

(3) *The mental content that guides motor action is not reflectively accessible.*

(4) *The mental content that guides motor action is not perfectly accurate.*

These claims are all supported by a wealth of empirical evidence, which I will outline for each. (1) and (2) might be dubbed "preliminary claims," as both feature only indirectly in subsequent chapters. Rather than relegating these preliminaries to the status of implicit assumptions, I have opted to discuss each at length, in the hopes of eliminating any confusion that might otherwise arise. Claims (3) and (4) assume a more primary status in this project, and will together constitute a crucial premise in the following chapter's argument against the factivity of knowledge.

This chapter will follow a simple format, with one section devoted to each claim in turn. This will include a detailed explanation of the claim and corresponding empirical support, which will include some combination of behavioural, neuropsychological, neurophysiological, and neuroimaging evidence. For the most part, relevant objections and philosophical accounts will be reserved for the next chapter, and understood specifically in terms of how they might impact my argument against factive knowledge. However, I will flag both particularly salient overlap with extant philosophical discussion and notably controversial empirical findings.

0. Reflective Access

Before discussing the above empirical claims, I first want to briefly clarify what I mean by "reflective access." Most crucially, I think framing reflective access in epistemic terms would be a mistake, especially for this specific project. That is, while reflective access can certainly result in knowledge, I will not take it to reduce to a form of knowing (see, for example, Littlejohn 2015). Rather than in epistemic terms, it is more useful to think of the reflective accessibility of mental content in relation to our capacity to report that content, incorporate it into our reasoning, and so forth. That is, I will be using the term in the sense that reflectively accessible content is, at least in normal circumstances, reportable content. This, I think, is more in line with the usage we might find in the empirical study of reflective access (see, for example, Higgins and Johnson 2012).

1. Preliminary: Unconscious Perception

There is a wealth of evidence suggesting that our cognitive systems often represent information despite us lacking all awareness of such representation. As my argument in chapter three will take this claim as something like a background assumption, I would like to take the time to discuss it now. This section will be devoted the phenomenon of perception without conscious awareness, "unconscious perception," and its empirical status. In short, while some dissent might be found, it is widely held that perception can occur without awareness, and multiple lines of empirical evidence support this conclusion.

To begin with, it is important to distinguish between distinct claims that might be appropriately labelled "unconscious perception." First, there is the claim that perception can decouple from phenomenology entirely, and that we might form perceptual states without any associated conscious experience. While this claim is well supported in its own right (see Matthen & Prinz 2015), this is unnecessarily strong for our purposes. The cases I focus on in the next chapter involve ordinary perceptual beliefs, and the information represented is certainly accompanied by conscious experience. Accordingly, I will use a weaker, less controversial sense of "unconscious perception," on which we might perceive things without conscious *awareness* (ibid, p. 372).[12] Further, if we take conscious awareness to be akin to something like reflective access,[13] this leads us to a more precise formulation of my first preliminary assumption:

> UNCONSCIOUS PERCEPTION: Not all perceptually acquired mental content is reflectively accessible.

Multiple lines of empirical evidence support the claim that perception can occur without reflective access, and there are three classes of evidence I will discuss: behavioural observations for individuals with neurological disorders (neuropsychological studies), behavioural observations for healthy individuals, and neuroimaging from healthy individuals.

First, there are specific neurological conditions that might result in perception without reflective access. As these unconscious perceptions can be much more conspicuous than those in healthy individuals, such cases are especially instructive. Here I will focus on the unconscious perceptions we might observe in the case of blindsight, in which "patients with an impaired visual cortex [can] perform visually in their blind field without acknowledging that performance" (Kroustallis 2005, 31). First, Pöppel observed that when patients with scotoma (regions of blindness in the visual field) were asked to guess the location of visual

[12] It is important to not forget that my account will be built on a weaker sense of unconscious perception than will be used by many of the researchers I will cite. While their stronger claim entails mine, it is open to criticism that mine is not. This sort of point will reappear in chapter three.
[13] This too appears in Matthen and Prinz (2015).

targets placed within those blind areas, and then indicate by looking in the direction of their guess, their responses correlated with the position of the targets (1973). Despite reporting no conscious perception of the target positions, the saccades (eye movements) of their guesses were not random, but systematically related to target position (ibid).

Even more remarkable are observations of the patient known as "DB" (for a complete case study see Weiskrantz 1986). In addition to displaying the same correlation between saccade and target location, DB was able to accurately reach for targets located in his scotoma, and even demonstrated "excellent performance" in discerning between "X" and "O" figures (Sanders et al. 1974, 708). It is worth repeating that from the perspective of DB, these attempts would have presented as nothing more than guesses, seemingly reaching at random into his blind spot. He had no conscious perceptions of the targets. As summarized by the researchers that first documented this behaviour (emphasis mine):

> The present results indicate that visual stimuli can be localised with considerable accuracy and that certain aspects of their orientation and spatial distribution can be differentiated if they are larger than a critical size when they are presented in a "blind" region. The patient's performance was systematically related to both the size and the duration of the stimuli although there ***was no acknowledged awareness of this capacity*** . . . Throughout these experiments he insisted that he saw nothing except in his intact field. (ibid, 708)

In short, blindsight offers a compelling case for unconscious perception. It has been observed in a multitude of subsequent studies (for a review of evidence and criticisms, see Cowey 2004), for tasks as complicated as navigating around a series of obstacles (de Gelder et al. 2008), all with subjects reporting no conscious awareness of their perception, or even conscious experience of any kind.[14] The key for our purposes is that the behaviour of individuals with blindsight indicates that they form perceptual representations to which they have no reflective access. While this provides a valuable illustration of what exactly perception without reflective access might look like, and seems to confirm that it does indeed occur, these extraordinary cases can only get us so far. Accordingly, I will now turn to evidence of unconscious perception in ordinary individuals.

The first set of observations involving unconscious perception (unrelated to neurological damage) I will discuss are those obtained under visual masking paradigms. Under such an experimental paradigm, "a clearly visible brief display can be rendered invisible (masked) when presented in close spatial and temporal proximity with other visual stimuli" (Davis & Kim 2011, 1990). Masking stimuli might be presented before the target stimuli ("forward masking"), after the target stimuli ("backwards masking"), or concurrently with it ("lateral masking") (see ibid; Wilkinson et al. 1997). In many such experiments, stimuli masked from awareness can nevertheless be observed to impact behaviour. It is this feature of visual masking that allows for compelling demonstrations of unconscious perception.

[14] It's actually more complicated than that, as researchers distinguish between blindsight with no conscious experience (type-1) and blindsight with some conscious experience (type-2) (see Kentridge 2015).

A typical masking study might employ a priming paradigm in order to demonstrate the behavioural influence of unconscious perceptions: After being briefly presented with a masked stimulus ("primed"), subjects are presented with a second "target" stimulus, the processing of which will vary with the representation of the masked stimulus (see Van den Bussche et al. 2009). For example, in one backwards-masking study (Alameda et al. 2003), participants were briefly (57 or 84 ms) shown a priming stimulus (in this case a word or phrase), followed by a masking stimulus, followed by a target stimuli comprised of Arabic numerals. The researchers observed that participants named target stimuli more quickly (i.e. "747") when they were preceded by an associated priming stimulus (i.e. "Boeing") (ibid). This sort of result indicates that the masked stimuli are still represented by the subjects. That is, the subjects' behaviour suggests perceptual representation of stimuli for which there is no reflective access. Studies confirming these types of results abound (for a meta-analysis see Van den Bussche et al. 2009).

Next, there is "crowding," an effect similar[15] to lateral masking, in which "a shape can be more difficult to identify when other shapes are near it" (Parkes et al. 2001, 739). For example, in one study, several Gabor patches were placed around a central patch, rendering the orientation of the central patch invisible to conscious perception when viewed peripherally (ibid). Remarkably, however, despite being unable to report the orientation of a single patch in the array, subjects were still able to accurately report the average orientation of the patches (ibid). This too indicates an unconscious perception of the orientation of the individual patches. While subjects clearly didn't have access to representations of individual orientation, the fact they could access the average orientation suggests that in these cases "the orientation signals in primary visual cortex are pooled before they reach consciousness" (ibid, 739).

Crowding effects have also been observed by utilizing priming. One such study elicited what is known as "orientation-selective adaptation" from crowded stimuli, an effect in which subjects exhibit quicker response times in identifying the orientation of a target stimulus when they are first primed with a stimulus of the same orientation as the target (Rajimehr et al. 2003). In this study, subjects displayed orientation-selective adaptation to crowed stimuli, "even when perceptual awareness of orientation had been confounded by crowding" (ibid, 1199). "Performance of subjects in reporting the orientation of crowded illusory lines was at chance level," confirming that there was no conscious perception of the orientation of the crowded stimuli (ibid, 1199). Again, this strongly suggests that subjects acquired perpetual representations of the orientation of the priming stimuli and simply lacked reflective access to those representations. This effect has also been observed for stimuli rendered (consciously) invisible by other means, such as "motion-induced blindness" (Montaser-Kouhsari et al. 2004).

In short, "unresolvable orientations substantially influence perception at multiple levels," (Rajimehr 2004, 663). Perhaps the most compelling illustration of this orientation-selective adaptation comes from a study by He and MacLeod, which didn't render orientation features consciously imperceptible by masking or

[15] Although experimentally quite similar to lateral masking, some researchers maintain that crowding effects are processed differently than masking effects (see Parks et al. 2001).

crowding, but simply by making them to fine to be consciously resolved (2001). Even when subjects were primed with stimuli of alternating black and white lines so fine that they were "perceptually indistinguishable from a uniform field," subjects still displayed adaptation to the orientation of these stimuli (ibid, 474). Again, these results support the claim that we might form perceptual states to which we lack reflective access:

> Our observations imply that extremely fine details, even those too fine to be seen, can penetrate the visual system as far as the cortex, where they are represented neurally without conscious awareness. (ibid, 474)

Another valuable illustration of unconscious perception comes from interocular suppression, in which "dissimilar images are presented to the left and right eyes, which leads to the suppression of one of the images from conscious perception" (Kaunitz et al. 2014; see Faivre et al. 2014 for a comparison with crowding). For example, one study found that "interocularly suppressed (invisible) erotic pictures" could still impact the attention of participants (Jiang et al. 2006, 17048). Another study found that a brightness illusion in which two identical circles appear to have different levels of brightness because of their surroundings persisted even when the background was suppressed beyond conscious perception (Harris et al. 2011). There is even evidence for high-level representation occurring without awareness, with one study finding that "emotional facial expression continues to be processed even under complete suppression" (Adams et al. 2010, 205). In short, as with the previous examples, interocular suppression presents compelling evidence that subjects form perceptual representations of visual stimuli they do not consciously perceive and cannot access.

The final type of behavioural evidence of unconscious perception I will discuss here comes from the phenomenon of inattentional blindness, "the failure to consciously perceive otherwise salient events when they are not attended" (Ward & Scholl 2015, 722). In one famous example of the effect, subjects were shown a complex video stimulus involving a number of people and asked to attend to a specific event (this amounted to some variation of counting basketballs being passed between specified individuals) (Simmons and Chabris 1999). During the video, either a woman holding an umbrella or dressed in a full gorilla costume would walk across the screen, but this highly unusual event failed to be consciously perceived by many participants (ibid). Depending on the configuration of the experiment, as few as 8% of those evaluated even noticed the gorilla, demonstrating that "without attention, we may not even [consciously] perceive objects" (ibid, 1059 & 1068). This effect has been demonstrated by numerous studies (Mack & Rock 1998; Lathrop et al. 2011; Seegmiller et al. 2011; Pitts et al. 2012; Drew et al. 2013; Schnuerch et al. 2016), and while I have no need to commit to the claim that conscious perception *cannot* occur without attention (see for example Mack & Rock 1998), it is at least clear that conscious perception often does not occur without attention.

Nevertheless, a solid chunk of evidence suggests that "stimuli that suffer [inattentional blindness] may be processed extensively by the perceptual system," and representations still form of stimuli hidden from conscious perception by inattentional blindness (ibid, 20). Mack and Rock describe a series of experiments that demonstrate this utilizing a priming paradigm (see ibid, chapter 8). As with the previous cases in which stimuli were suppressed from conscious perception by other means, stimuli rendered consciously imperceptible via inattentional blindness also result in priming effects, demonstrating they are indeed perceived. Again, this is supported by the findings of multiple other studies (Koivisto & Revonsuo 2007; Lathrop et al. 2011). Additionally, there is also even evidence for semantic processing of hidden stimuli. That is, the "meaning of objects undetected due to inattentional blindness interferes with the classification of attended stimuli" (Schnuerch et al. 2016, 459). Again, this result is further evidence for perception without reflective access.

The behavioural observations presented up to this point might be summarized in the following way:

> There is now ample evidence in the literature that sensitive, indirect methods of testing often reveal that perceptions not consciously experienced seem to be encoded, and facilitate or inhibit subsequent perception when that same or a related stimulus object is subsequently presented to the observer. (Mack and Rock 1998, 173)

While this sort of behavioural evidence is most valuable for our purposes, there is an independent line of neurological evidence for unconscious perception, which I want to close out the section by at least mentioning. "Representation of stimulus orientation" can be detected by fMRI, even when that stimulus is rendered consciously invisible by forward and backwards masking (Haynes and Rees 2005, 868). To be clear, this means that orientation of an invisible stimulus unconsciously perceived by a subject can be accurately predicted from fMRI of that subject (ibid). "Amygdala activation" has been observed to correspond with face stimuli that were interocularly suppressed (Williams et al. 2004, 2898; see also Pasley 2004). Finally, stimuli hidden by inattentional blindness also appear to produce an observable neural signal (Pitts et al. 2012). In short, for much of the behavioural evidence suggesting that consciously invisible stimuli are still represented, there is corresponding neurological evidence that confirms this.

In summary, a wealth of evidence strongly supports the claim that perception can be unconscious, particularly on the weaker reading of the claim on which some perceptual states are merely inaccessible. This idea, that we have no reason to assume reflective access to all our perceptual content, is crucial to keep in mind going forward. If you go in operating under the assumption that all our perceptual content is are accessible, it's unclear how any subsequent sections or the chapters beyond might be made sense of.

2. Preliminary: Multitudinous Representation

It is natural to imagine that perceptually acquired information, such as the location of an object, is represented internally exactly once, corresponding with a single mental state. So conceptualized, this view, "supported by subjective experience, is that we construct a single spatial map of the world in which objects and actions are represented in a unitary framework" (Colby & Goldberg 1999, 320). However, if we look beyond the subjective experience of perception, we find compelling evidence that this naïve conception is incorrect. Perceptual information is not represented singularly, but multitudinously:

> The space around us is represented not once but many times in parietal cortex. These multiple representations encode locations and objects of interest in several egocentric reference frames. Stimulus representations are transformed from the coordinates of receptor surfaces, such as the retina or the cochlea, into the coordinates of effectors, such as the eye, head, or hand. The transformation is accomplished by dynamic updating of spatial representations in conjunction with voluntary movements. This direct sensory-to-motor coordinate transformation obviates the need for a single representation of space in environmental coordinates. (ibid, 319)

In this section, I will discuss this multitudinous representation, focusing specifically on the representation of perceptually acquired spatial information used in motor action.[16] Accordingly, the claim I intend to support might be thought of as something along these lines:

> MULTITUDINOUS SPATIAL REPRESENTATION: Perceptually acquired spatial information is represented multiple times by our cognitive systems, in multiple frames of reference.

Before presenting a selection of the empirical evidence in support of this claim, I should provide a bit of disambiguation by elaborating on different sorts of spatial representations, certain features they exhibit, and the relationship between them. The core idea here is that "processes in the sensorimotor chain can be described in terms of internal representations of the target or limb positions and coordinate transformations between different internal reference frames" (McIntyre et al. 2000, 2823).

The first representations we form of perceptual stimuli are coded in the reference frame of the relevant sensory receptor: "Visual stimuli are initially coded in retinal coordinates, tactile stimuli relative to the skin surface, and auditory stimuli relative to the head" (Pritchett et al. 2012, 437). As visual perception is what I am most interested in for the purposes of this and subsequent chapters, here I will focus on initial representation of visual stimuli in the reference frame of the retina, "retinocentric representation." With every saccade (movement of the eye between fixations), "any location information encoded in the retinocentric reference frame that is predominant in the visual system is obliterated" (Schneegans & Schöner 2012, 89). Thus, these low-level representations are constantly updating.

Through a series of coordinate transformations, these retinocentric representations are transformed into other, more stable representations. The

[16] I have chosen to focus on spatial information not because this is where evidence for multitudinous representation ends, but for clarity and simplicity. As my subsequent argument deals largely with ordinary beliefs comprised of perceptually acquired spatial information, going beyond this wouldn't be to the argument's benefit.

details of this process are still the subject of current research.[17] However, this isn't that important for our purposes since here it's only important to recognize that we don't form unitary representations. On a basic model, retinocentric signals from each eye are first combined into a binocular retinocentric representation, which is then transformed into representations in egocentric (body-centered) and finally allocentric (world or object -centered) (Wade and Swanston 1996). Egocentric reference frames might include shoulder-centered (Flanders et al. 1992), eye-centered, or hand-centered coordinates (Chen et al. 2014). Conversely, allocentric representations are those representations "of the spatial location of an object relative to that of another external object, independent of the ego's position or orientation, whether present, imagined, or remembered" (Filimon 2015, 2).

Additionally, we might note that (i) representations from different sensory modalities need to be transformed into common reference frames before being useful, (ii) a process that contributes to motor variability—both points that are important for later discussion:

> Using multiple sensory modalities for planning similar movements is potentially problematic, as different sensory signals arrive in different reference frames. Specifically, early visual pathways represent stimulus location relative to current gaze location (a retinotopic representation), whereas proprioceptive signals represent hand location relative to the shoulder or trunk (a body-centered representation). To utilize these signals, some of them must be transformed between reference frames. Although sensory transformations may appear mathematically simple, we found that transformations can incur a cost by adding bias and variability into the transformed signal. (Mcguire & Sabes 2009, 1056)

There are multiple lines of evidence that support this general picture of multitudinous representation, and here I will look at three: neuropsychological, behavioural, and neuroimaging.

Much as with the example of blindsight in the previous section, here too I will begin with evidence for multitudinous representation that emerges from a specific neurological disorder, this time spatial hemineglect. However, unlike in the previous section, this disorder doesn't simply offer evidence that those with hemineglect exhibit multitudinous representation, but that we all do. Operating under the core assumption of cognitive neuropsychology, that lesions don't fundamentally alter the architecture of the brain but merely disengage specific cognitive systems (discussed in the previous chapter), spatial hemineglect offers a valuable insight into ordinary cognition.

When individuals suffer lesions to a hemisphere of the brain, they often present with a limited ability to "explore the side of space contralateral to the lesion (contralesional), and to report stimuli presented in that portion of space" (Vallar 1998, 87). Interestingly, however, if lesions are unilateral (present in only one

[17] Even some of the basic claims in this section are contested. For example, it has been argued that there are actually no allocentric representations, only egocentric (Filimon 2015), that transformations can occur directly from retinocentric to allocentric (Souman et al. 2010), and that multiple reference frames are used simultaneously in movement planning (McGuire and Sabes 2009). Again, the important point is only that we represent spatial information in multiple ways.

hemisphere of the brain) these limitations are often isolated to contralesional space, with "patients' performance . . . comparatively preserved in the side ipsilateral to the lesion (ipsilesional)" (ibid, 87). When patients with unilateral lesions experience diminished cognitive capacities in contralesional space, it is known as "spatial hemineglect." For our purposes, spatial hemineglect can be extremely informative, as cognitive capacities often breakdown in systematic ways, with damage often impacting "specific sectors of space with reference to a given coordinate system" (ibid, 87). That is to say, as spatial neglect often impacts a patient's ability to form representations in certain reference frames, spatial hemineglect provides compelling evidence that these spatial representations are indeed distinct, with lesions selectively disrupting certain reference frames.

Consider the following example from Driver et al. (1994): Subjects with left spatial neglect, resulting from damage to the right hemisphere of the brain, were shown a stimulus comprised of a line of equilateral triangles aligned at the base. For certain iterations, a small gap would be present on either the left or right side of the middle triangle. However, the stimulus was always presented at one of two possible diagonal orientations, so that the gap would always appear at the same location on the screen. Remarkably, despite this, the participants demonstrated a much higher failure rate for identifying gaps in the triangle when they were on the left. This indicates the stimuli were represented in an object-centred reference frame that Driver et al. called "axis-based coordinates," in reference to a vertical axis one might imagine running vertically through the centre triangle (ibid, 1362). After all, the location of the gap was "the same in every possible reference-frame except the axis-based coordinates," on which left neglect can then explain the results (ibid, 1362). In short, these observations provide a dramatic illustration of the powerful effects exerted by the axis-based reference-frame chosen during shape perception" (ibid, 1362). Similar experiments have found specific stimuli represented in multiple other reference frames, even varying with distance from the body (see Humphreys et al. 2013; Li et al. 2014).

Next, I will discuss behavioral evidence for multitudinous spatial representation. Studies that evaluate pointing errors can be especially valuable for this purpose, as we might expect different error patterns depending on whether spatial information is represented singularly or multitudinously. Here, I will first provide a brief account of the approach as a whole, followed by a specific example of one such study, and a summary of the conclusions we might make with respect to multitudinous representation.

To begin with, these behavioral studies typically involve pointing at a virtual target or a target that is only displayed before the act of pointing begins. The idea here is that this eliminates the possibility of "movement corrections based on tactile feedback," which would otherwise contaminate observation of "the underlying frames of reference putatively used by the brain to specify the endpoint" (Battaglia-Mayer et al. 2003, 1009). Next, these studies exploit the fact that error in representation of spatial information would manifest differently depending on the reference frame. Consider a toy example (modified from McIntyre et al. 2000, §3): Say we have a cognitive system with two possibilities for (2D) spatial representation: a singular reference frame of Cartesian coordinates or a singular reference frame of polar coordinates. As both random error (noise) and constant systematic error (bias) will manifest differently in Cartesian vs. polar coordinates, there are certain patterns of pointing error we might expect for spatial information represented in each. For example, in the case of multiple pointing attempts at a given target over the course of a study, trials stacking up along a radially outward trajectory will be indicative of inward bias ("overshoot") in polar coordinates. While human processing of spatial information is of course much more complicated than a binary choice between two coordinate systems, the same principle still holds. Certain error patterns can be indicative of specific sequences of spatial representation and coordinate transformation, plus specific sources of error.

One such study conducted by McIntyre et al. analysed the error for subjects pointing at targets (which were presented and removed before the pointing phase) in two lighting conditions: (1) a "lighted condition" in which subjects could barely see their fingers, and (2) a "dark condition" in which subjects couldn't see their fingers at all (ibid, 2835). For the lighted condition, the error pattern indicated "a view-centered reference frame for the representation of the final pointing position" (ibid, 2839). The dark condition resulted in a far more complex error pattern that "at first glance do not point to a clearly defined reference frame" (ibid, 2841). However, subsequent simulations showed that at least a portion of the error pattern could be explained by "the cascading effects of two transformations, one centred at the eyes and the second at the shoulder" (ibid, 2851). That is, pointing error in the dark condition suggests the utilization of three distinct spatial representations with two coordinate transformations in between.

These findings by McIntyre et al. are by no means unique, and they themselves highlight "data from three different experiments support the hypothesis of a two-stage transformation" (ibid, 2851; citing Carrozzo et al.1999; Baud-Bovy & Viviani 1998). Similarly, Flanders et al. concluded that, for these specific pointing tasks, "the neural representation of target parameters is transformed from a retinocentric representation to a shoulder centered representation" (1992, 309). More generally, the analysis of point error has produced evidence that, given test conditions, (1) "specification of the final position of the hand occurs in a viewer-centered reference frame", (2) "eye and hand information combine into a common egocentric binocular frame of reference," and (3) "irrespective of the specific frame of reference used to plan the movement, the different spatial parameters of reaching are not treated in a unitary manner, but are processed in parallel and largely independent of each other" (Battaglia-Mayer et al. 2003, 1010). The specifics themselves aren't as relevant for our purposes as what they collectively demonstrate: multitudinous spatial representation.

Beyond the neuropsychological and behavioral, direct observation of the brain has also yielded a wealth of evidence for multitudinous spatial representation. Here I will briefly discuss how neuroimaging studies support the claim that spatial information is represented in multiple ways. These sorts of studies typically involving using fMRI to monitor brain activity while subjects perform specific tasks, which will vary depending on the target representations (for reviews see Galati et al. 2010; Boccia et al. 2014). For example, a study by Zaehle et al. (2007), which investigated differences in the neural substrates for egocentric vs. allocentric representations, asked participants questions that were designed to elicit either a much stronger egocentric or allocentric response. A typical allocentric question might begin, "The blue triangle is to the left of the green square. The green square is above the yellow triangle. The yellow triangle is to the right of the red circle," and then ask, "Is the blue triangle above the red circle?" (ibid, 99). An egocentric query, on the other hand, would define position in strictly egocentric terms. In observing the corresponding neural activity, the researchers concluded that "data show separate neural circuits mediating different spatial coding strategies" (ibid, 91). Similarly, other studies found that "allocentric and egocentric reach mechanisms use partially overlapping but different cortical substrates" and "different networks of areas subtend egocentric and allocentric strategies" (Chen et al. 2014, 12515; Boccia et al. 2014, 236). Beyond a simple egocentric/allocentric division, neuroimaging studies also suggest that "body referencing, i.e., the integration of spatial body knowledge . . . into visual and somatosensory maps" is distinct from other egocentric representations, and that "long-term cognitive representation of stable spatial features of a familiar environment" distinct from allocentric representation generally (Galati 2010). In short, neuroimaging studies don't only provide compelling evidence for multitudinous spatial representation, but they support the same representational divisions as independent behavioural and neuropsychological studies. This further strengthens the claim that spatial information is represented multiple times, in multiple ways.

In conclusion, there is a strong empirical case for multitudinous representation. This claim, and particularly the egocentric/allocentric distinction, foreshadows elements of the discussion in the following section, which examines observations demonstrating that the representations of conscious perception are not shared by motor action. Here we might note that this result is not at all unexpected, as successfully acting on objects might require representations in egocentric coordinates, whereas making perceptual judgements about objects might require representations in allocentric coordinates (see Milner and Goodale 2008, §3-4). For example, if I want to reach out and grasp my coffee mug, the ideal coordinate system for representing the location of the mug is something egocentric, like hand-centered. Conversely, if I want to judge whether that mug is left or right of the coffee maker, the ideal coordinate system for the mug is something allocentric, like coffee-maker-centered. The core idea here is that visual judgement and visually guided motor action require representations with distinct properties, which, as we'll see shortly, is central to my core argument against factivity.

3. Motor Action

In the previous two sections, I introduced what I have considered preliminary claims. While the argument in chapter four won't explicitly hinge on them, and any references to them will be sparse, one might struggle to understand my project without them in the background. Especially as they might be somewhat unfamiliar to an epistemological audience, it was important to provide these background assumptions a bit of visibility. With these in mind, I will now introduce the empirical claims that will together form a key premise in the next chapter's argument against the factivity of knowledge. This section will focus on the mental content that guides motor action, discussing the empirical evidence that the representations used for action are distinct from those of conscious perception. The next section will focus on noise and bias in that same content.

In their seminal 1992 paper, Goodale and Milner proposed that there are separate visual pathways for conscious perception and action. That is, "the set of objects descriptions that permit identification and recognition may be computed independently from the set of descriptions that allow an observer to shape the hand appropriately to pick up an object" (20). They based this proposal on both neuropsychological and neurophysiological[18] evidence for independent ventral (recognition, "what") and dorsal (action, "how") streams of visual processing, noting that the two systems have different "output requirements," needing representations with different properties to successfully function (ibid, 20). Over the past three decades, multiple lines of evidence have converged on the idea that "visual control of action is distinct from [conscious] visual perception" (Westwood and Goodale 2011, p. 804)[19]. In this section, I will draw heavily from this research programme to support the following claim:

> INACCESSIBLE MOTOR GUIDANCE: The mental content that guides motor action is reflectively inaccessible.

Before detailing the evidence for this claim, let me first provide a bit of a methodological disclaimer. While the question of whether conscious perception guides motor action has certainly been debated within philosophy (see Clark 2001, 2007, 2009 for the affirmative; Wallhagen 2007 & Mole 2009 for the negative), this section will strategically ignore this literature. The reason for this is simply that I wish to construct an argument against factivity that rests directly on empirical claims, committing myself only to those philosophical assumptions that are implicit in the empirical research. Of course, philosophical objections to the empirical claim presented in this section are still highly relevant, and I will return to them in the next chapter, when discussing objections that one might raise against my account. Keeping this in mind, I will now turn to the empirical case that motor action is guided by inaccessible content, discussing in turn evidence from neuropsychology, behavioural studies, and neuroimaging.

The neuropsychological evidence that led Goodale and Milner (1992) to posit distinct visual pathways for conscious perception and action centres around two distinct neurological conditions, optic ataxia and visual agnosia. Both conditions result from lesions to the brain, and each constitutes one half of a double dissociation between representations for conscious perception and representations

---

[18] These were studies in which monkeys had electrodes surgically implanted into their brains, which were then monitored the activity of specific neurons as the monkeys performed pre-trained tasks (see Caminiti et al. 1990 for a detailed account of standard methods). Due to the horrifying ethical implications of such experimentation, I have opted to exclude them from my own account.

[19] Note that much of the literature pertaining to this vision-for-perception/vision-for-action distinction uses "perception" in a narrow sense, meaning something approximating "perceptual experience." As I will use a broader sense of perception, closer to "gaining information via perceptual capacities," I will clarify the use of the narrow sense by specifying *conscious* perception (see Milner and Goodale 2008, 775). This is important to recognize, as it allows us to understand vision-for-action as a sort of unconscious perception.

for motor action: Patients with optic ataxia display all the behaviours associated with ordinary conscious perception, but struggle to perform successful motor actions. On the other hand, those with visual agnosia display behaviours that indicate limited (conscious) perceptual capacities, but this doesn't present a problem for successful motor action. Taken together, this suggests that the two visual systems are largely distinct, as one can be damaged without that damage impacting the other.

Let's begin with optic ataxia. As already mentioned, the condition results from lesions to the brain (for description of lesion sites see Perenin & Vighetto 1988, 646-7). At the behavioural level, patients with optic ataxia struggle to successfully reach for visual targets ("misreaching") (McIntosh 2010). For some this can occur even when visual targets are viewed directly, but more often misreaching only occurs for peripheral targets (ibid). For example, one such patient "performed large reach errors with the left hand to peripheral targets presented both in the left and right visual fields, and with the right hand mainly in the left visual field." (Ferrari-Toniolo et al. 2014, 60). However, what makes the condition rather remarkable is that despite consistent misreaching, patients displaying optic ataxia often "have intact visual fields, stereoscopic vision, oculomotor control, proprioception, motor abilities, and cerebellar function" (Andersen et al. 2014, 968). Despite being able to provide an accurate "verbal estimate of the distance or the relative position of objects," patients nevertheless struggle to successfully act on those objects (Perenin & Vighetto 1988, 648). For Milner and Goodale, as visual perception is intact while motor action is impaired, this suggested a dissociation between the two. For our purposes, the crucial bit is that patients with optic ataxia don't seem to lack any of the reflectively accessible perceptual content we might expect from a "normal" individual. As deficits are "independent of visual space misperception" (Perenin & Vighetto 1988, 643), i.e. conscious perceptual experience is unaffected, as are motor abilities, misreaching in optic ataxia provides evidence that reflectively inaccessible representations guide motor action (which lesions can adversely effect).

Next, visual agnosia provides evidence that motor action is guided by reflectively inaccessible spatial representations, via the converse dissociation. Here I will focus on the famous example of the patient known as "D.F.," who developed multiple (conscious) perceptual deficits after suffering "diffuse brain damage as the result of the carbon monoxide poisoning" (Milner et al. 1991, 421).[20] Milner and Goodale summarized her case in the following way:

> Despite her profound inability to recognize the size, shape and orientation of visual objects, DF showed strikingly accurate guidance of hand and finger movements directed at the very same objects. Thus, when she was presented with a pair of rectangular blocks of the same or different dimensions, she was unable to distinguish between them. When she was asked to indicate the width of a single block by means of her index finger and thumb, her matches bore no relationship to the dimensions of the object and showed considerable trial to trial variability. However, when she was asked simply to reach out and pick up the block, the aperture between her index finger and thumb changed systematically

[20] For a more recent case study of a different individual, see Karnath et al. (2009).

with the width of the object, just as in normal subjects. In other words, DF scaled her grip to the dimensions of the objects she was about to pick up, even though she appeared to be unable to 'perceive' those dimensions. (1992, 22)

Most remarkable for our purposes is that D.F.'s lack of reflectively accessible spatial representations did not impede her ability to successfully act on them. D.F. displayed severe visual deficits in the conscious perception of motion and depth in addition to orientation (see Milner et al. 1991), which led to her spatial judgements, such as the "orientation of a large open slot," to be "severely impaired" (ibid, 417). However, her motor action still appears as if guided by accurate representations of spatial information, suggesting it is guided by perceptual content "not accessible to her perceptual awareness" (ibid, 418). In short, as she seemed to lose certain reflectively accessible spatial representations without a corresponding loss in successful motor action, we can understand D.F.'s case as clear evidence that the spatial representations that guide motor action are reflectively inaccessible.

Moving beyond lesion studies, we also find evidence for reflectively inaccessible guidance of motor action in the response of "normal" individuals to certain visual illusions. Namely, while somewhat controversial, a number of illusion studies seem to demonstrate dissociation between conscious perception and action. Key for our purposes is that visual illusions can impact reflectively accessible spatial content without impacting corresponding motor action, suggesting motor action is guided by inaccessible content immune to the illusion. Aglioti et al. first described dissociation in this context in their seminal 1995 paper, and Haffenden et al. produced similar findings (1998). However, as methodological objections have been voiced against these studies (Franz et al. 2000, 2008; for defence see Westwood & Goodale 2011, §3.2), here I will discuss an example study that replicates the results while adhering to the procedural recommendation of Franz et al. (2003).

In a study conducted by Ganel et al., participants were observed reaching for cylindrical objects placed on an illusion-inducing surface (2008). The objects, which differed in length, were placed on the surface in such a way that lengths would be judged incorrectly, and unsurprisingly, "participants erroneously perceived the physically longer object as the shorter one (or vice versa)" (ibid, 221). Nevertheless, their motor actions were unaffected by the illusion: "their grasping was remarkably tuned to the real size difference between the objects" (ibid, 221). In short, these results indicate that "action and perception can be fully dissociated when real and illusory object size are pitted against" (ibid, 223). For our purposes, this provides further evidence for the claim that motor actions are guided by reflectively inaccessible spatial information: Despite the accessible spatial content, on which participants formed their judgements, representing object size incorrectly, the content that guided their motor actions represented the size correctly. Therefore, it appears that inaccessible spatial content guided the motor action unaffected by the illusion.

Finally, neuroimaging studies offer compelling support for the claim that the mental content guiding motor action is reflectively inaccessible. This follows from the wealth of evidence indicating that the neural substrates of conscious perception are distinct from those of motor action. First, returning to the example of D.F., results of fMRI performed on her "confirm directly that D.F.'s visual form agnosia is associated with extensive damage to the ventral stream, and that her spared visuomotor skills are associated with visual processing in the dorsal stream" (James et al. 2003, 2463). More broadly, evidence from neuroimaging provides an independent case that the mental content that guides motor action is reflectively inaccessible, because, as argued by Kozuch (2015, see §4), such evidence suggests no representations are both conscious and capable of directly guiding motor action ("efficacious").

Kozuch first points out that, "while the ventral stream probably produces conscious representations, these representations appear to not feed directly into motor action" (ibid, 582). Beyond studies involving lesion cases, he highlights evidence that motor action doesn't appear to increase activity in ventral stream areas of the brain. For example, one such study by Culham et al. (2003), using fMRI to observe ventral and dorsal activity during motor action, found that while reaching-and-grasping tasks resulted in greater dorsal activity than reaching alone, ventral activity did not increase. From this, they concluded, "These results suggest that dorsal areas . . . but not ventral areas . . . play a fundamental role in computing object properties during grasping" (ibid, 180). Conversely, Kozuch notes that, while "there is no reason to doubt that the dorsal stream produces representations that feed directly into motor action," it does not appear that such representations are conscious (2015, 585). Returning to lesion studies, the primary evidence he cites for this comes from fMRI of patients with optic ataxia. In these cases, lesions to specific dorsal areas responsible for guiding motor action "fail to bring about deficits in visual consciousness" (ibid, 588). This, of course, strongly suggests that these representations that guide motor action do not contribute to conscious perception.

This account is strengthened by the findings of a neuroimaging study by Cavina-Pratesi et al. (2007), which observed activity in ventral and dorsal areas for tasks such as "grasping, reaching, size discrimination, pattern discrimination, [and] passive viewing" (1). They found that dorsal activity displayed no modulation for perceptual tasks not involving motor action, whereas ventral activity did (ibid). Conversely, ventral activity didn't modulate with motor tasks, while dorsal activity did (ibid). In short, "these areas were differentially activated depending on whether the task was grasping or perceptual discrimination" (ibid, 1). They concluded:

> The human visual system does not construct a single representation of the world for both visual perception and the visual control of action. Instead, areas in the ventral stream mediate the visual perception of objects whereas areas in the dorsal stream mediate the visual control of action directed at those same objects. (ibid, 12)

To summarize, a wealth of empirical evidence from neuropsychology, behaviour, and neuroimaging together create a compelling case for Milner and Goodale's two visual systems hypothesis. Most relevant for this project is how this helps demonstrate that the mental content that guides motor action is reflectively inaccessible. In a broad sense, this follows from the inaccessibility of information represented by the dorsal stream. However, even if we don't want to fully commit to two visual systems, we might understand this evidence as still supporting the inaccessibility claim introduced in this section, as I've been careful to frame this discussion to illustrate. In doing so, I hope to have shown that there is a strong empirical case that the mental content that guides motor action is reflectively inaccessible.

4. Visuomotor Error

In the previous section, I introduced a claim that will form one half of a bipartite empirical assumption in the next chapter's argument against factive knowledge. In this section, I will discuss the second half of this assumption, the comparatively simpler claim that the mental content that guides ordinary motor action is not perfectly accurate. Of course, there is a trivial sense in which this claim is correct, as sometimes unsuccessful motor action might be attributable to inaccurate mental content: We might say an unsuccessful catch was caused by "misjudging" the trajectory of the object in flight, and so forth.[21] This is not what I have in mind. Rather, my claim is that even successful motor action can be guided by mental content that is not perfectly accurate, and moreover that we have compelling empirical reason to suppose this is typically the case. In short, this claim can be formulated in the following way:

> VISUOMOTOR ERROR: The mental content that guides successful motor action is not perfectly accurate.

The empirical support for this claim assumes a fairly simple structure. The mental content that guides motor action will inevitably contain some degree of systemic error. Sometimes this error will be significant enough that the corresponding motor representation cannot guide successful action. However, much more relevant for our purposes is that most of the time this error will not reach levels that it precludes successful action. Such error might lead to slight variability in reach vectors or grip aperture and might be observable in a controlled setting, especially in settings where visual feedback is supressed, preventing mid-action corrections. However, this error is beyond our awareness when we perform motor actions. In the next chapter, I will argue for the epistemological significance of this fine-grained error. Here I will only introduce the empirical case for it, which involves two types: Random error (noise) and systematic error (bias).[22]

Let's begin with noise. Upon even the most cursory reflection, we might recognize that motor action is unavoidably, "universally, sometimes frustratingly, variable" (Churchland et al. 2006a, 1085). This variability is most obvious for tasks like throwing, for which even miniscule variances can impact the success of an action. However, even for more mundane tasks such as reaching or grasping, the variability inevitable for motor action is unmistakable. While of course some

[21] Here we might note that, given the considerations from the previous section, "misjudging" might not be the best term.

[22] To be clear, this is not to say that all noise might be considered random error (or all bias systematic error), but something like the converse.

variability might be attributed to environmental factors, it nevertheless persists even when "external conditions, such as the sensory input or task goal, are kept as constant as possible" (Faisal et al. 2008, 292). Movement variability that is independent of external conditions can be attributed to noise, "random or unpredictable fluctuations and disturbances that are not part of a signal," acquired at various visuomotor stages (ibid, 292). For example, certain noise is associated with the "movement execution stage," on which "planned motor commands are sent to the muscles so that the movement is actually made" (van Beers et al. 2004, 1050). Known as "execution noise," this can be attributed to the physiology of muscle fibers themselves, (see Faisal et al. 2008), and is therefore acquired independently of the mental content that guides motor action. Execution noise, however, is not the only internal source of variability in motor action, and there is compelling evidence for noise in the mental content itself.

Much of the variability we might observe in motor action is attributable to noise that is acquired before movement execution begins. This is "planning noise," visuomotor noise that "arises from noisy sensor estimates of hand and target position and during the coordinate transformations required" to specify a particular movement (Apker et al. 2010, 2655).[23] While it is unclear the exact proportion of motor variability that can be explained by planning noise, both behavioural (see ibid) and human neurophysiological (see Churchland et al. 2006a/b) evidence suggests it is a primary contributor to variation in ordinary reaching tasks. In short, it is clear that planning noise plays a non-trivial role in the variability of motor action. This, I maintain, is sufficient to conclude that the mental content guiding motor action is not perfectly accurate. In order to illustrate why, it might be helpful to first look at an example study.

In an experiment designed specifically to investigate the interaction between planning and execution noise, Apker et al. observed subjects as they reached for various virtual targets projected in three-dimensional space, tracking how their movements varied depending on both target location and whether their hand was visible during the task (2010). While details of their analysis are likely both too technical and unnecessary for the purposes of my project, the basic idea is this: We would expect different error patterns in reaching tasks depending on the source of the error (execution or planning noise). Thus, by observing the actual patterns in reaching error, one can determine the relative contributions of execution and planning noise. Given the observed distribution of reaching error, this particular study concluded that "visual planning–related noise plays a dominant role in determining anisotropic[24] patterns of endpoint variability in three-dimensional space" (ibid, 2654). Specifically, errors in these reaching tasks were "elongated along the depth axis" in such a way that they indicated "increased planning noise associated with localizing targets in depth" (ibid, 2662 & 2664). Interestingly, this evidence even suggests that the visual feedback from seeing one's hand during motor action "does not reduce sensory uncertainty in depth" (ibid, 2661). As Apker et al. point out, this is consistent with the findings of

[23] Note that while here I will be take "planning noise" to include "sensory noise" (noise involving initial target and hand position representation), the two are often discussed as distinct (see Shi and Buneo 2012).

[24] Distributed differently along different dimensions

independent research into vision suggesting inaccuracy in depth localization (ibid, 2662; citing McIntyre et al. 1998; van Beers et al. 1998, 2002; Viguier et al. 2001). This, when taken together with other studies that focus specifically on noise (Churchland et al. 2006a/b; Apker and Buneo 2012; Shi and Buneo 2009, 2012; Osborne et al. 2005) constitute a strong empirical case for the ubiquity of planning noise.

I think it is clear that the best way to understand this noise, present in both initial sensory information and subsequent representation of spatial information in successive coordinate systems, is as inaccuracy, however minor, in the mental content that guides motor action. First, planning noise affects the representation of spatial information and subsequent motor commands used to guide motor action, which is just to say it affects the mental content that guides motor action. While noise in motor commands isn't directly relevant for our purposes, noise in the representation of the spatial information that precedes these commands is. This noise, like that in the representation of the depth of visual targets described by Apker et al. (2010), might be easily understood as a sort of inaccuracy. The noise in representation of the depth of a target in 3D visual field is nothing more than representational deviation from the actual position of the target. Our representations of target position deviate from the actual position. This is inherited by our motor commands and ultimately results in slight motor error. Put simply, the mental content that guides motor action is not perfectly accurate, slightly misrepresenting the spatial properties of objects under action.

In case it isn't clear that this noise should be understood a sort of error or inaccuracy, the slight misrepresentation of spatial information, let's consider what seems to me the only obvious alternative: inexactness. Rather than less accurate representations, one might instead say that planning noise results in less determinate representations of spatial information. Perhaps we don't slightly misrepresent quantities such as position and size, but instead inexact ranges, correlating with noise levels. While this might seem tempting, there are two reasons to resist this. First, as argued by Hellie (2005), it seems intuitive that we at least represent colours maximally determinately, but that doesn't preclude noise in the perception of colour. This doesn't make sense if noise is understood as increasing indeterminacy. Moreover, there is a larger problem for positing inexactness. If we say that motor representations are inexact to the point that this inexactness can explain motor variability attributable to planning noise, then it's unclear how these representations might be used to explain any specific motor action. If a representation of, say, an object's location is inexact to the point that reaching for either of two distinct points A or B is consistent with this representation, how do we explain the fact that an individual reached for point A during trial 1 and point B during trial 2? How can the same inexact representation explain two distinct actions? The entire purpose of motor representations is to explain motor action, something they can't do if we say they're so inexact that they're consistent with any number of distinct actions. Clearly, if we frame noise as inaccuracy, we can explain such observations. For trial 1, the subject represented the object as being located at A, and so forth. However, if the representations are inexact, we can no longer appeal to them to explain specific action at a desired level of determinacy.

Ultimately, a compelling body of evidence points to the role of planning noise in motor variability. For ordinary motor actions, planning noise levels won't be high enough to impede our successful motor action or even for us to notice. However, as any noise in the representation of the spatial information used to guide action amounts to inaccuracy, a non-zero difference between, for example, represented object location and actual object location, we can understand a clear sense in which the mental content guiding motor action is simply not perfectly accurate.

Beyond error attributable to noise, there is also compelling empirical evidence for systematic bias in the mental content that guides motor action. As I will discuss in more detail in chapter five (§2), it has been widely observed that representations of spatial information are biased towards our center of vision (Müsseler et al. 1999; Sheth and Shimojo 2001; Kerzel 2002; Adam et al. 2008; Fortenbaugh & Robertson 2011).[25] This "foveal[26] bias," on which participants "systematically [mislocalize] the target closer to the center of gaze" (Sheth and Shimojo 2001, 329), is often studied using perceptual judgements (see Odegaard et al. 2015). However, there is also evidence for foveal bias in visual representations that guide motor action. For an example of this, let's look at a series of experiments performed by Bock, on which participants pointed at projected targets, with movement restricted to an arc in the horizontal plane (1986). As participants' hands were occluded from their view, when they pointed in the direction of a target, they had no visual feedback to correct for initial misrepresentation of target location; thus, any such misrepresentation became "manifest as pointing errors" (ibid, 476)[27]. Bock's results from these experiments first confirmed "a tendency to point consistently past the targets by a constant amount, which was independent of target position for a given subject, but varied among subjects" (ibid, 477). He concludes that the collective effect of pointing past the target "suggests that the eccentricity[28] of visual targets on the peripheral retina is over-estimated," specifically observing an average of 3.83° of overestimation of target location radially (ibid, 481). Similar results have been observed by other studies (Prablanc et al. 1979; Lewald 1997; Henriques et al. 1998; Ross et al. 2015).

Now, of course, this bias is much more pronounced in the experimental conditions than we might expect for everyday action, as we normally can make use of visual feedback during the course of motor tasks. However, the key here is that while visual feedback might minimize the effects of foveal bias, it does not eliminate it entirely (see Brenner et al. 2008).[29] In this way, we might understand

[25] Sheth and Shimojo (2001, 329) note that these sorts of observations date to at least the nineteenth century (citing Helmholtz 1866).

[26] The fovea is the center of vision on the retina.

[27] It is worth noting, as Bock points out, that pointing error in these studies can also be caused by "deficits of motor control," not just misrepresentation of target position (ibid, 476). Therefore, these sorts of studies must and do take care to "study visual localization in paradigms where no time constraints are imposed upon the movements, and where pointing accuracy is independent of movement amplitude," in order to minimize the contributions of deficits in motor control (ibid, 477).

[28] The angle on the retina from the fovea

[29] Additionally, some might worry that this bias is specifically in retinocentric representations, and motor action is guided by representations transformed from retinocentric coordinates. However, as coordinate transformations generally tend to *increase* error, not eliminate it (Schlicht and Schrater

that all the ordinary representations of locations in space that guide motor action will be slightly biased towards the centre of visual space. Along with planning noise, this bias amounts to an additional sense in which the mental content that guides successful motor action is not perfectly accurate.

5. Conclusion

With this chapter, I hope to have provided sufficient empirical groundwork for my subsequent argument against factive knowledge. As will become apparent shortly, the critical claim here is the inaccuracy of the mental content that guides motor action. Without inaccuracy of mental content, and therefore false beliefs, there is no argument against factivity. However, it is important to not lose sight of the role played by claim (3), that this content is not accessible upon reflection. In order for my argument to be even remotely plausible, it must be the case that we cannot access the erroneous content, given that we don't seem to notice it. In the parallel arguments against factivity I present in chapter five, this inaccessibility is more complicated, as we have some reflective access to content but not the error of that content. However, everything is simplified greatly if we can just observe, as with motor action, that the content is simply inaccessible to begin with.

2007), we have no reason to think such misrepresentation would not persist in the content that guides motor action. Indeed, Bock (1986) posited such misrepresentation specifically because of observation of motor error.

Chapter Three: Against the Factivity of Knowledge

In this chapter, I will present an argument in favour of replacing factivity with some weaker constraint on knowledge. I have opted to formulate this argument as a series of initially plausible assumptions from which an undesirable conclusion naturally follows. The factivity of knowledge is one of these assumptions, and the primary empirical claims discussed at length in the previous chapter will form another. Beyond this, we need only assume our ordinary practice of belief attribution to reach a conclusion inconsistent with our ordinary practice of knowledge attribution. I will argue at length that factivity is the easiest of these assumptions to do away with.

In so arguing against factive knowledge, this chapter will assume a simple structure. In the first section, I will introduce the three plausible assumptions and show how an undesirable conclusion directly follows. In the sections following, I will anticipate a variety of objections, as well as responses to this putative problem other than my preferred move of replacing factivity: In section two, I will argue that we have multiple reasons to retain the empirical claims advanced in chapter two, most compelling of which is an independent conceptual argument to the same end. In section three, I enumerate the exact cost of discarding either our ordinary practice of belief or knowledge attribution, and anticipate certain objections that might arise at this stage, most significant of which is the idea that the mental content specified in assumption (1) is simply not the same type of content of interest to epistemologists and/or implicit in assumption (2). Section four will be devoted to the assumption that knowledge is factive, and I intend to demonstrate that replacing it with a weaker constraint comes at a modest epistemological cost, and is therefore the most attractive response. In the final two sections, I address additional concerns one might have with this argument. Section five addresses the objection that empirical evidence suggests that ordinary knowledge-attributing practices are often unreliable, and section six discusses concerns with the argument's blend of naturalistic and folk epistemology.

1. The Basic Argument

This section and the rest of this chapter concern three initially plausible assumptions, the reasons we might have for retaining or rejecting these assumptions, and the relationship between them. These assumptions are as follows:

(1) The mental content that guides motor action is (a) reflectively inaccessible and (b) not perfectly accurate.

(2) Perceptual beliefs guide motor action.

(3) Knowledge is factive.

There isn't much to say on the initial plausibility of these assumptions. The entire preceding chapter was devoted to establishing the expansive empirical case for (1). Recall that we find compelling neuroimaging (James et al. 2003; Culham et al. 2003), neuropsychological (Westwood and Goodale 2011), and ordinary behavioural evidence (Aglioti et al. 1995; Ganel et al. 2008) for (1a) stemming from Milner and Goodale's broader claim that the visual guidance of motor action is distinct from conscious perception (Goodale and Miler 1992; Milner and Goodale 1995; 2008). Similarly, a wealth of empirical evidence suggests the mental content that guides motor action displays both random error (noise) (Osborne et al. 2005; Apker & Buneo 2012) and constant error (bias) (Bock 1986; Ross et al. 2015). While this error doesn't ordinarily preclude successful action, it is nevertheless clear we cannot consider such content to be perfectly accurate, therefore establishing the initial plausibility of (1b). While I will return to the matter in section three, the initially plausibility for (2) is quite simple to explain, as it follows directly from our ordinary practice of belief attribution. Namely, we are perfectly willing to attribute a perceptual belief to an agent on the basis of her motor action alone. For example, we might attribute to S the belief that an object was "there[30]" on the exclusive basis of her attempt to act on that object, successful or otherwise.[31] Of course, assuming the reliability of ordinary mental state attribution might be immediately thought to be in some sort of tension with the naturalistic approach on which we establish assumption (1), and I will return to this sort of concern in sections five and six. Until then, even if they appear to emerge from opposite poles of the methodological spectrum, which they do, we can at least agree on the initial plausibility of both. Finally, the initial plausibility of assumption (3) is simply a matter of it being all but universally accepted in epistemology, as discussed in chapter one. Beyond that, there is little to say on the matter. Regardless of whether we ultimately wish to retain all three assumptions, it should be clear each is at least initially plausible.

The problem with committing to these three assumptions is that, as I will argue, it leads to following undesirable conclusion:

(C) The everyday perceptual beliefs that guide successful motor action cannot be known.

Before explaining why the above assumptions lead to this conclusion, let me first say a bit on what makes this conclusion so undesirable. To put it simply, accepting such a conclusion requires us to go against our ordinary practice of knowledge attribution. When an agent performs a successful motor action, we readily attribute all sorts of perceptual knowledge to her. Distinguishing it from incidental contact

[30] To be clear, this isn't the belief it is "there" in the sense that there *is* an object, but rather a specific belief about the location of that object.

[31] It may be worth noting that this holds even for cases we wouldn't attribute such a belief on the basis of S's gaze orientation. If an object isn't in her direct line of sight, and instead on the periphery of her visual field, we withhold any attribution of a belief regarding the object's position until she attempts to act upon it, at which time we attribute to her the belief that that object is there.

and involuntary movement, successful object-directed action requires[32] an agent knows the object's location and size, and most likely its orientation. We have no problem attributing knowledge of, for instance, object location on the basis of successful action, and our intuitive judgements will invariably be that she knew that the object she successfully acted upon was there. Why was she able to successfully reach out in the direction of the coffee cup? Because she knew it was there, which is to say her belief that the cup was there counts as knowledge. Why was she able to successfully grasp the handle? Because she knew it was oriented towards her, which is to say that her belief that the cup was oriented towards her counts as knowledge. Whatever instance of successful motor action we might consider, it seems obvious that the beliefs that guide those actions are usually known. We cannot retain this ordinary practice of knowledge attribution if we accept (C). Instead, we are faced with a skeptical result, as a non-trivial swath of our putative knowledge would actually be unknown. At the very least, such a result is undesirable.

To understand why the above three assumptions lead to this conclusion, let's start with (1) and (2). Assumption (1) specifies two properties of the mental content that guides motor action, and (2) specifies that belief states guide motor action. That is, (2) says, or might be understood to say, that whatever mental content guides motor action is belief content. Accordingly, (1) specifies two properties of the belief states that guide motor action. Per (1a), the beliefs that guide motor action are reflectively inaccessible, and per (1b), those beliefs aren't perfectly accurate. At this point, one might interject that such a result already seems an undesirable conclusion, as we don't want to deny reflective access to perceptual beliefs that guide successful motor action. This discussion I will leave for section three. All I will say here is that it's important to not confuse such a claim with the more troubling one that the perceptual beliefs that we seem to obviously have reflective access to regarding the location, size, and orientation of objects are actually inaccessible. In keeping with Goodale and Milner's description of the two visual systems, I am not saying that we don't actually have access to such beliefs. Rather, I am saying that that accessible belief content is not the belief content that guides motor action. Beyond our accessible perceptual belief content, there is inaccessible content that guides motor action, and, of course, is also not perfectly accurate.[33]

Next, if we grant assumption (3) that knowledge is factive, it follows that beliefs described by (1) and (2) cannot be known. In order to see this, let's consider an object at location (x, y, z), which is represented as being at location $(x + \delta_x, y + \delta_y, z + \delta_z)$.[34] As discussed in the previous chapter, the noise and bias inevitable in

[32] As will be discussed in section four, this is a bit hasty. However, even if we don't want to go so far as to say these beliefs must be known, we can at least agree that these beliefs are often known. The same problem follows from this.

[33] As we'll see in chapter five, this picture is further complicated by the potential for inaccessible noise in otherwise accessible content.

[34] Two notes here: First, as discussed in the previous chapter, it is likely that this sort of spatial information might be represented in spherical, not Cartesian coordinates. However, which reference frame or coordinates are actually used doesn't matter for this point. Additionally, while objects obviously take up regions in space, not singular points, this doesn't actually matter either. We could

these beliefs results in fine-grained inaccuracies in the representation of spatial information, such as object location. That is to say, δ_i values will be non-zero, and objects will be represented as being somewhere that they are not. Under a factive conception of knowledge, beliefs with such content cannot be known. If a belief represents an object as being located somewhere it is not, regardless of the magnitude of this misrepresentation, that belief is false.

Of course, we need to proceed with caution. One might object that this is an unrepresentatively strong take on factivity, and that we often allow for a certain level of approximation for known beliefs. 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.[35] 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.[36] Accordingly, it seems that our tolerance for approximation doesn't extend to the position of objects as represented in the content of belief. While we may have no problem attributing knowledge in cases where the *sentences* used to describe beliefs are only approximately true, this doesn't appear to hold when the belief contents themselves are only approximately true.[37]

We can further illustrate this by considering that what it means for knowledge to be factive is simply that S knowing that p entails p. Imagine we allow for approximate truth in factive knowledge. Then for a case in which S knows that p, where p is approximately true, from factivity this entails p. However, because p only approximates the truth, there will be some proposition p*, which describes the way the world actually is, for which p ∧ p* is a contradiction. Accordingly, we might

---

run the same argument with 3D regions in space, but the math involved in describing the space occupied by even the simplest shapes is prohibitively complex for our purposes.

[35] I'm playing it a bit fast-and-loose here with the properties of mental content, relying only on hypothetical belief reports and not something more robust. Nevertheless, I think this is acceptable for the purposes of illustration.

[36] We might imagine she has any amount of justification (etc.) for her belief. If the cup isn't exactly at the edge, we still judge that her belief is false and therefore unknown.

[37] These types of cases pose their own interesting problem, as it seems we're willing to attribute knowledge for the approximately true, inaccessible beliefs that guide motor action, but not approximately true beliefs like the one that the cup is exactly on the edge of the table. For now I will simply say this problem doesn't impact the latter's ability to evidence our intolerance for approximate truth with respect to the factivity of knowledge. I will return to this in section one of chapter four.

understand that if knowledge is factive, beliefs that are only approximately true cannot be known, regardless of how close to truth they might be.

In this manner, we can understand why the beliefs described by (1) and (2) cannot be known under (3). However, even when faced with the mere approximate truth of the everyday perceptual beliefs that guide motor action, I would submit that our intuitive judgement is that such beliefs ought to count as knowledge. In section three, I will return to the question of whether we should indeed consider such beliefs to be known. Until then, it should at least be clear that this result is certainly undesirable. Yet it is the direct consequence of three assumptions, for which the rejection of any one appears undesirable in its own right. In the following sections I will argue that rejecting factivity is the easiest solution to this problem, by a wide margin. In the process of arguing for this, I will also address objections as they arise, with the intention of strengthening the claim introduced in this section that these three assumptions indeed pose a problem that cannot be solved without either the outright rejection of at least one or the begrudging acceptance of (C).

2. Why Retain the Empirical Assumption?

Setting aside concerns about whether the problem described in the previous section is a real problem, the most pressing of which I will return to in section three, dropping (1), the empirical assumption that the mental content that guides motor is both reflectively inaccessible and not perfectly accurate, likely seems the most attractive solution. After all, assumption (1) describes a sort of mental content that might be somewhat unfamiliar from the perspective of mainstream epistemology. While here I will avoid the more difficult question of whether epistemologists even ought to be concerned with such mental content, again circling back to this in the next section, as a descriptive matter, it is clear that inaccessible, perceptually acquired content with fine-grained error properties is outside the standard epistemological wheelhouse. Moreover, while certainly supported by a wealth of evidence, it's not as if (1) enjoys the same degree of empirical support as, say, the claim that the speed of light in vacuum is constant. In this light, simply disregarding assumption (1) likely seems an appealing option.

In this section, I will present four reasons for resisting the impulse to disregard (1). The first is simply that the primary philosophical arguments that have been levied against the broader project of Milner and Goodale, one by way of an objection to Clark (2001), don't impact the narrower claim of (1a), and there doesn't appear to be any clear, extant philosophical objection relevant to (1b). Second and more generally, I think the impulse to disregard (1) on the basis of its empirical nature is fundamentally misguided. However, even if we discount the epistemological value of such empirical claims, we still have multiple compelling conceptual reasons to commit to (1). Not only might a counterfactual do similar theoretical work to the empirical evidence detailed in the previous chapter, albeit not as convincingly, but we might additionally construct independent conceptual arguments for both (1a) and (1b). In doing so, I hope to demonstrate that rejecting (1) is not the easy out it might appear.

As mentioned in the previous chapter, beginning with his 2001 paper, Clark has used the Goodale and Milner program to argue for a claim similar to (1a). Clark first notes that philosophers generally commit to the following "Assumption of Experience-Based Control (EBC)":

> Conscious visual experience presents the world to the subject in a richly textured way, a way that presents fine detail (detail that may, perhaps, exceed our conceptual or propositional grasp) and that is, in virtue of this richness, especially apt for, and typically utilized in, the control and guidance of fine-tuned, real-world activity. (ibid, 496)

However, Clark argues that the Goodale and Milner program provides compelling evidence that this is mistaken, and instead offers an alternative, the following "Hypothesis of Experience-Based Selection (EBS)":

> Conscious visual experience presents the world to a subject in a form appropriate for the reason-and-memory-based selection of actions. (ibid, 512)

The crucial omission from EBS is of course motor control: While conscious visual experience guides the selection of action, what I'll discuss below in the context of "providing reason for action," it doesn't guide action itself. As the claim that EBC is mistaken and conscious visual experience isn't an apt guide for motor action is quite close to (1a), I clearly need to say something about objections to Clark's account. Here I will look at Wallhagen (2007), after which I will consider Mole's (2009) objection directed at Milner and Goodale. In both cases, I will argue that neither of these objections threatens (1a).

Let's being with Wallhagen's objection to Clark. Most relevant for our purposes is Wallhagen's account of D.F.'s visual agnosia, the foundational evidence for Milner and Goodale discussed at length in the previous chapter. Wallhagen argues that Clark is mistaken in concluding that D.F.'s case suggests that EBC is incorrect, as it is plausible that D.F. actually retained her capacity for conscious visual experience of "the shapes, sizes and orientations of things" (2007, 556). Instead, Wallhagen contends that D.F. simply is unable to "form a perceptual judgment about aspects of form," which doesn't itself entail that D.F. lacks a conscious visual experience of the same (ibid, 556). Additionally, while not in direct response to the illusion studies discussed in the previous chapter as evidence for (1a), Wallhagen makes similar claims for visual illusions. Specifically, Wallhagen contends that while our conscious perceptions represent visual illusions accurately, "subjects simply do not notice that their visual states have this content" (ibid, 554).

Regardless of the extent to which this this threatens Clark's account or is even correct, these objections pose no threat to (1a) for the simple reason that Clark's claim is far stronger than (1a). For the purposes of my argument against factivity, we don't need to assume that the mental content guiding motor action is independent of conscious visual experience, although that would certainly get the job done. Rather, all that is necessary is that such content is inaccessible. Wallhagen does nothing to suggest the content guiding D.F.'s motor action is accessible, and indeed concedes that D.F. cannot "perceptually identify certain aspects of perceptual stimuli, to bring those features under concepts" (ibid, 556), which is to say, *access* those features. The same holds for Wallhagen's account of visual illusions. Even if we grant that subjects in the action-under-illusion experiments (chapter two, section three) consciously experience the correct size of the objects they act on, this doesn't suggest that they have reflective access to the correct size of those objects. As with D.F., their perceptual judgements indicate they do not. In short, because (1a) is a claim only about reflective access, not conscious visual experience, it is not subject to Wallhagen's objections to Clark, which only target the latter.

If we turn to Mole (2009), we find something similar. Mole's objections are pointed directly at Milner and Goodale's claim of two distinct visual systems, and he argues that both action-under-illusion and visual agnosia can be understood in terms of a single visual system. First, Mole argues that results from illusion studies can be explained in terms of contradictory conscious representations, without needing to posit separate visual pathways for each. In this case, the perceptual information such as object size that guides motor action is "represented through embodied demonstratives" (Mole 2009, 1002). For example, when a subject thinks "'This disc is this big' . . . the content of the 'this' is specified by a bodily gesture, such as the grip aperture of a grasping hand" (ibid, 1002), which can account for why such contradictory content isn't obvious. Mole argues similarly in the case of D.F. As he maintains that "the content of an experience can be given through an embodied demonstrative" (ibid, 1008), Mole concludes that we can understand D.F.'s gestures as providing the phenomenal content that guides action, thus eliminating the need to posit a separate, unconscious visual system guiding action. However, again, even if we grant that Mole is entirely correct here, this account does not threaten assumption (1a) as it seems implausible to suggest that subjects in illusion studies or D.F. have reflective access to the content specified by embodied demonstratives. After all, in both illusion studies (Ganel et al. 2008) and descriptions of D.F. (Milner et al. 1991) there are no indications that perceptual reports change after action occurs. Moreover, it's not immediately clear how we might ever upon reflection gain access to the spatial information represented by these embodied demonstratives. As with Wallhagen's objections to Clark, even if we grant Mole's claim that embodied demonstratives provide the content that guides motor action, it appears such content is still reflectively inaccessible.

Next, I think we should resist any impulse to disregard (1) simply on the basis that it is an empirical claim. For my project specifically, this relates back to the naturalistic method discussed in chapter one. Since one of my primary aims is to increase the points of contact between epistemology and relevant empirical sciences, I certainly don't consider disregarding difficult empirical evidence to be a serious option. Beyond this, as a more general matter, I cannot imagine that there is a legitimate philosophical rationale behind ignoring strong, relevant empirical evidence. Of course, one might argue that against specific empirical evidence on the basis that it is either irrelevant or too weak, but this is not at all the same thing as disregarding it. As the previous chapter was devoted to arguing for the strength of this evidence in great detail, and this chapter might be understood as arguing for its relevance, there does not appear to be any good reason to ignore the empirical case for (1).

Furthermore, even if we disregard assumption (1) in its capacity as an empirically derived descriptive claim about actual mental content, we might take inspiration from this claim to construct a counterfactual that does similar theoretical work. Consider some case in which an evil demon or scientist or whomever alters someone's mind so that the mental content that guides her motor actions is both reflectively inaccessible and not perfectly accurate in the precise way as described in the previous chapter. We can then run the same sort of argument as provided in the previous section, but now specifically for her beliefs. My thought is that we still want our theories of knowledge to be capable of handling these beliefs, from which the argument can continue as if we had actually accepted (1). Of course one might object that we don't actually need our theories of knowledge to handle these beliefs, but as far as I can tell, this is analogous to the objection that beliefs with content described by (1) are not of epistemological interest, with the same arguments for and against (which will be discussed in the next section). Nevertheless, I do recognize that this type of argument doesn't seem as convincing. Moreover, we might also note that this route doesn't allow for an anti-skeptical stopgap argument, which requires that our actual beliefs have the properties described by (1).

Finally, I would maintain that we have compelling conceptual reason to posit (1) entirely independent of any empirical considerations. Let's begin with (1a). If we consider any object-directed motor action, the precision required for that action to be successful seems to confound the mental content we can access upon reflection: While any accessible content seems too imprecise to guide successful action, and is just as compatible with failure, any content precise enough to guide an action to success seems implausible to posit as accessible. Take the example of catching a baseball in flight. Of course, before, during, and after the catch, there is a host of mental content that we might report on.[38] We might report the ball is "above us" or "about 20 feet up" or "arching towards us." At the moment of the catch, we might report the ball was "right above us" or "directly in front of us." However, none of this is precise enough to credit with guiding with successful action, and is just as consistent with placing our hand in the region of the ball but failing to catch it. Moreover, when we try to dial up the precision of this content so that it might be credited with guiding a successful catch, it is no longer plausible to say such content is reflectively accessible. Whatever the precise location of the ball at the moment of contact with our hand, we do not have reflective access to it. However, this doesn't preclude our ability preform a successful catch. That is to say, the mental content guiding that catch is, at least to some degree, reflectively inaccessible. Of course, this isn't due to any feature unique to catching a ball. Any time we reach for an object, it is implausible to say that we have reflective access to the precise location and size information necessary to result in consistently successful action, and not a series of flailing, grasping motions in the general direction of that object. In short, on account of nothing more that the discrepancy between the precision required for successful motor action and imprecision of our reflectively accessible mental content, we have compelling conceptual reason to

[38] Recall from the previous chapter that the capacity to report on content is key to the sense in which I'm using "reflective access."

conclude with (1a) that the mental content that guides successful motor action is reflectively inaccessible.

Furthermore, I think we also have compelling conceptual reason for positing (1b). Or, more precisely, I think something like (1b) seems overwhelmingly intuitive. Were we capable of directly observing the mental content that guides motor action, it would be a remarkable surprise to learn that we represented object location with perfect accuracy. Of course, without appeal to empirical evidence it is difficult to say conclusively that we might commit fine-grain motor error attributable to misrepresentation. However, this too seems overwhelmingly plausible. If we grant this, it follows that we cannot appeal to inexactness as an alternative to inaccuracy, for, as argued in the previous section, such inexact representations are unable to explain how a particular misrepresentation might result in a particular motor action. In this manner, we might understand the conceptual case for (1b) as a simply matter of it seeming extremely likely. Perhaps this isn't independent reason to accept (1b), like we have for (1a). However, I do think it makes it that much more difficult to dismiss (1b) offhand. The empirical evidence confirms what we strongly suspected.

In the end, for the multiple reasons presented in this section, as well as the significant empirical case presented in the previous chapter, it would be a mistake to reflexively reject assumption (1). Beyond this, such a move comes at the cost of divorcing knowledge from empirical descriptions of mental content, on the basis of preserving a constraint on knowledge. I will return to this matter in section four, and argue that it is fundamentally mistaken to commit to this cost on such a basis.

3. Why Retain Ordinary Belief and Knowledge Attribution?

Given that you are inclined to grant assumption (1), either on the basis of the empirical evidence, conceptual argument, or counterexamples that might do comparable theoretical work to (1), it is still unlikely that you are willing to reject the factivity of knowledge. Instead, I think it is most likely you might object to some part of the claim that our ordinary belief and knowledge attribution practices pose any real problem for factivity, even given that we assume (1). Moreover, even if you do agree with my argument that (1)-(3) leads to (C), it is likelier you consider the cost of revising our belief and/or knowledge -attributing practices lower than that of rejecting factivity. In this section, I intend to address both of these potential responses, using my answers to obvious objections to both clarify the costs of rejecting (2) or accepting (C), and strengthen the argument outlined in section two. While different approaches might encounter different problems, every one comes at the cost of divorcing knowledge and action.

This section will assume the following structure: To begin, I will clarify what content I have in mind for the perceptual beliefs that guide motor action. After this, I will consider three possible ways one might attempt to block my argument at this point. First, one might reject assumption (2) outright, maintaining that I am mistaken in taking beliefs to be the mental states that guide motor action. Next, I consider what appears to be the strongest objection to my account, which argues that the beliefs I described are not of epistemological interest. Finally, one might also simply opt to accept (C) outright. Beyond the shared cost of divorcing knowledge and action, each of these responses encounters specific problems, which I will also discuss in detail.

As discussed in section one, the implications of accepting (1) and (2) seem simple enough: (1) specifies certain properties of the mental states that guide motor action, but doesn't specify which states those are. (2) specifies that beliefs are the mental states that guide motor action, but doesn't specify any properties of belief states (beyond, of course, whatever properties we implicitly take to be fundamental for belief states). Therefore, if we grant (1) and (2), we can conclude that the belief states that guide motor action have content with the specified properties of the content that guides motor action: The content of these beliefs is not reflectively inaccessible, and they are not perfectly accurate.

This move is reminiscent of that from Butterfill and Sinigaglia, who maintain that the "contents of intentions can be partially determined by the contents of motor representations" (2012, 119). Now there is some disagreement on what exactly motor representations are (see Ferretti 2016). However, the key here is that motor representations in some way "represent action outcomes" (Butterfill and Sinigaglia 2012, 120). For example, when I reach for an object, that reach is guided in part by my motor representation of that reaching motion, in addition to my representation of the location of that object. Accordingly, Butterfill and Sinigaglia are making a parallel move here, arguing that specific content that guides motor action needs to be counted in the content of the intentional states. While they aren't concerned with belief states specifically, their motivation is broadly the same as mine. If we want to say that any given mental state explains action, the content of that mental state needs to include those representations that guide or "coordinate" action (ibid, 119).

Let's consider an example of the content of a perceptual belief regarding an object's location under object-directed action. Say, under everyday conditions, S reaches to grasp object O. If we freeze this example mid-reach, we have no problem attributing a belief of O's location to S, and would likely express this as, "S believes that O is there." Part of the content of S's belief will be given in some manner by her conscious perceptual experience, that reflectively accessible content she might report, use in her reasoning, and cite as a reason for her action. However, as I have argued at length, this reflectively accessible content is insufficient to guide her ordinarily successful actions. In order for this belief to explain her action, we must also say that part of the belief's content is determined by that inaccessible, slightly inaccurate content that guides her action. In this manner, we might understand the content of her belief to be determined by both accessible and inaccessible content.[39] Of course, there is nothing special about object location, and the same argument works for size and orientation as well.

[39] One might object here and say that actually there are two beliefs, one corresponding to the ventral stream and the other to the dorsal. I think we should resist this move for the reason that this would result in multiple beliefs described by the same sentence, which refer to the same features of the same objects and jointly explain behavior. For the sake of simplicity, I want to say this is just one belief with content determined by both ventral and dorsal streams. Moreover, such a move doesn't avoid the cost of divorcing belief and action.

One benefit of this move is that it allows us to continue to cite the same beliefs as both guiding motor action and *providing reason for* motor action. As a matter of ordinary practice, we commonly treat beliefs like those that an object is "there" as simultaneously guiding action and providing reason for action. In an example of reaching for a coffee mug, not only would we say that the belief regarding the location of that mug, S's belief that the mug is there, but we might also cite this belief as providing reason for her action. In addition to her desire to drink coffee, etc., one reason she acted on the mug is that she believed that it was there. Given that the mental content that guides motor action is reflectively inaccessible, this feature is only preserved if we take both it and reflectively accessible perceptual content to together determine the content of the perceptual beliefs in question.

The first way one might object to this picture, while still accepting (1), is by maintaining that the mental content that guides motor action does not determine the content of belief states, but rather some other type intentional state.[40] Perceptual states seem like the best candidate for this move. Motor action could then be explained in terms of these perceptual states, whose contents have the properties specified by (1), without any unseemly epistemological consequences. However, there are at least two problems with such an account, the first of which is the cost. As this amounts to just denying (2), it comes at the expense of our ordinary practice of belief attribution: If beliefs don't guide motor action, then our practice of attributing beliefs on the basis of motor action would seem to be flawed. Motor action could never be explained in terms of belief, and more generally, any behaviour would be unexplainable in terms of belief to the extent that the behaviour consisted of motor action. Here I will go no further than simply describing these costs. In the next section I will return to them, and argue that such consequences represent a much higher cost than rejecting factivity.

Moreover, this response has another problem in that it only works if the properties described by (1) are unique to the mental content that guides motor action. If, for example, it turned out that all our perceptually acquired mental content was also subject to noise, or any noise-like fine-grained error, we would no longer be able to relegate such content to non-belief states.[41] As it happens, we have good reason to think that something like this might be the case. First, the intuitive plausibility of (1b), outlined in the previous section, generalizes to all visual perception, not just that used for guiding motor action, and of course we might construct a counterexample in which all perceptual content is subjected to fine-grained misrepresentation. Moreover, as empirical evidence suggests, "the very different neural processes of [conscious] perception and action are limited by the same sources of noise" (Osborne et al. 2005, 412). While still not reflectively

[40] One might motivate this by maintaining that when we attribute beliefs on the basis of motor action, we have very different content in mind than that described by (1). Namely, implicit in those belief attributions is an assumption that these beliefs have reflectively accessible content determined exclusively by the believer's conscious perceptual experiences. If it turns out that this assumption is incorrect, and that there are no beliefs with than content capable of actually guiding motor action, then we should withdraw these attributions.

[41] Of course we *could*, but this would eliminate the category of perceptual belief entirely. Within the setting of epistemology, this consequence is so obviously undesirable I won't consider it a reasonable option.

accessible[42], noise in conscious perception can be observed in the variability of our perceptual judgements (see Fiehler et al. 2010) and, as mentioned in the previous chapter, foveal bias is widely documented in perceptual judgements (see Odegaard et al. 2015). I will return to this matter in detail in section two of chapter five. As other objections can avoid this problem without incurring any additional cost, this doesn't appear the best route to take.

I will now turn to a second objection to my argument, which improves upon simply denying that beliefs guide motor action. One might grant that beliefs, with content as described by (1), do indeed guide motor action, but maintain that these aren't the sorts of beliefs epistemology is concerned with. The most common response I've encountered to this non-factive project runs along these lines, specifically maintaining that the beliefs in question are too fine-grained to be of epistemological interest. As I understand the objection, there are two ways it might be formulated: Epistemology might be uninterested in extant beliefs with fine-grained properties, or it might be uninterested in fine-grained descriptions of extant beliefs. On the former, there are just some beliefs that have fine-grained properties, whereas others do not, with epistemology only concerned with those without fine-grained properties. On the latter, we might describe any belief in either coarse or fine -grained terms, with epistemology only concerned with coarse-grained descriptions. It is important to make this distinction because the associated cost is different for the two formulations. For the first formulation, as the beliefs that guide motor action would have fine-grained properties alleged to put them beyond the realm of epistemology, (C) is an inevitable conclusion. As with the previous objection, this would serve to divorce belief from motor action. On the other hand, the second interpretation would allow for course-grained descriptions of those beliefs that guide motor action, effectively filtering out their fine-grained properties and allowing for (1)-(3) without accepting (C). Accordingly, this appears to be the more promising of the two routes.

Regardless of which formulation you take, I think this objection is fundamentally mistaken in attributing this ostensible extra-epistemological nature to grain. Simply put, it is unclear to me in what sense the belief content guiding motor action could be said to be excessively fine-grained beyond epistemological relevance. I suspect what is meant here is just that the scale of spatial quantities I'm using is marginally finer than some might be used to, resulting in the individuation of beliefs that otherwise might be identical.[43] For example, say I successfully reach for and grasp the same object twice. On the first attempt I represent the object's position as exactly where it is, but for the second attempt I represent it as one micron radially outward from its actual position. I want to say that these beliefs have different content; under factivity only the first counts as knowledge; therefore we should replace factivity with a weaker constraint that counts both as knowledge. The objection is that epistemology isn't interested in individuation that is this spatially determinate. Rather, on an epistemological account, these beliefs are considered identical. My response to this objection is that its account of

[42] Note that Hellie (2005) suggests, at least in the case of color, that this noise is actually somewhat accessible upon careful reflection, although this seems questionable.

[43] I don't think this objection can be understood in terms the content in question being non-propositional or non-conceptual, because, as far as I can tell, my account isn't committed to either.

epistemological boundaries is simply incorrect. It mistakenly takes a problem with inaccessible content to be one with fine-grained content.

We can understand the mistake in chalking the problem up to grain if we return to beliefs with "exact" content, like those discussed in section one, and filling in some details. Let's say someone believes that a cup is on the exact edge of a table, with not a single Planck length (ℓ_P) of offset or overhang, as is revealed by probing her belief contents. Perhaps she is able to report on the contents of her belief in terms of a Planck length. All that matters is that she believes the cup is on the exact edge. Now lets say the cup is actually $2\ell_P$ removed from being on the exact edge. Now first I think our intuitive judgement in such a case is that this belief is then false. It does not have the same truth conditions as the belief that the cup is $2\ell_P$ off the edge. Next, if asserted out of something like stubbornness, we would say the belief is unjustified. However, let's say instead our believer has an extremely powerful microscope that can resolve distances down to the Planck length, and observes that the cup is on the exact edge, but immediately after looking away a freak earthquake moves the cup by $2\ell_P$. We would now judge that she has a justified false belief. Moreover, if we instead say the cup is on the exact edge, we might construct a Gettier-style case where a calibration error of $2\ell_P$ is offset by the freak earthquake moving the cup $2\ell_P$ just as she looks away, giving her a justified true belief we don't judge to be known. The point in all this is that all our familiar judgements and epistemological framework seem to be preserved, even if we individuate beliefs down to the Planck length. If epistemology isn't supposed to work with beliefs simply because they involve fine-grained spatial descriptions, we certainly need an account of why it sure looks like it does.

Rather than epistemology not extending to fine-grained beliefs, I think it is much more plausible that a problem involving access is being mistaken for a problem involving grain. That is, the correct objection is that epistemology isn't concerned with beliefs with reflectively inaccessible content. After all, while the seemingly benign Planck length example involves reflectively accessible content, the beliefs I'm concerned with do not. We cannot report on their content nor use them in our reasoning. This I think is the most difficult objection for my account to handle, with the most difficult formulation given in terms of content description: Epistemology isn't concerned with descriptions of belief content on which that content is not reflectively accessible, and therefore we shouldn't require our theories of knowledge to handle such descriptions.[44] However, while I will concede that one consequence of my account is that some of the time some belief content will be inaccessible, I don't think this is as disastrous a result as it might appear. In order to see this, it's important to recall what beliefs I'm claiming will have partially inaccessible content.

As discussed above, my claim isn't that under object-directed motor action, for some motor representation of objection O being at (x,y,z), we might say that S unconsciously believes that O is at (x,y,z). While I don't think this belief attribution is necessarily misguided, it does seem that if S has no access to this content, there might be a case that these aren't the sorts of beliefs that epistemology is interested

[44] After all, at least in the case that a theory of knowledge results in this "content unawareness," it might be considered a major problem (Abath 2012).

in. However, we might instead capture the mental content that guides motor action (and provides reason for motor action) with demonstrative beliefs. Taking inspiration from both Mole (2009) and Butterfill & Sinigaglia (2012), we might say that S believes that O is there, with both reflectively accessible content from perceptual experience and reflectively inaccessible motor representations filling in the content. We might also do this for size (S believes that O is that big.) and orientation (S believes that O is in that orientation.).

In doing this, there are two key points I want to highlight. First, it certainly seems that a theory of knowledge should be capable of handling such cases. To begin with, we might note that these beliefs seem to align with specific knowledge-wh attributions:

LOCATION (K-WH): S knows where O is.

SIZE (K-WH): S knows what the size of O is.

ORIENTATION (K-WH): S knows what the orientation of O is.

Far from being epistemologically uninteresting, these types of knowledge attributions have been subject of a good deal of contemporary epistemological discussion (see Parent 2014). Moreover, consistent with the fact that (at least) most knowledge-wh seems to reduce to knowledge-that (see ibid), we might express this knowledge in an explicitly propositional format:

LOCATION (K-THAT): S knows that O is there.

SIZE (K-THAT): S knows that O is that big.

ORIENTATION (K-THAT): S knows that O is in that orientation.

I think it is perfectly reasonable that we should expect our theories of knowledge to be capable of handling these types of knowledge attributions. At the very least, I can think of no obvious reason why we should exclude such attributions from the analysis of knowledge.

Second, I want to highlight that, as we saw with Mole (2009), we don't seem to have a problem with the content of demonstrative beliefs being somewhat reflectively inaccessible. In most everyday contexts, when S believes that object O is "there," it seems unreasonable for us to expect S to be capable of reporting much of anything on the content of this belief, especially anything precise and quantitative that corresponds with the spatial information represented by her visual system. Of course, this doesn't mean that excluding reflectively inaccessible content and/or demonstrative belief attribution from our theory of knowledge is not an available option to preserve factive knowledge, but only that the beliefs described by (1) and (2) aren't obviously beyond the domain of epistemology. Accordingly, any such argument must overcome the problem of explaining why these seemingly benign beliefs are actually extra-epistemological.

Nevertheless, one might still hold that, all things considered, we *shouldn't* require our theories of knowledge to handle such beliefs. As with simply rejecting (2), it seems we must leave this open as an option. However, this move does come with certain costs. First, as with the previous objection, this will still serve to divorce epistemology from motor action. As discussed at length, the mental content that guides motor action will always be reflectively inaccessible, regardless of how much we pixelate our description. Moreover, as with rejecting assumption (1), this position has the added cost of distancing the epistemologically viable descriptions of belief content from empirically informed descriptions of mental content. I will return to both these problems in section four, and argue that they represent a higher cost than is associated with including this content and accordingly replacing factivity with something weaker.

To close out this section, I would like to discuss one final route one might take in objecting to my account, simply accepting the conclusion (C) that those beliefs that guide our motor actions are unknown. Perhaps if the beliefs that guide motor action have the properties specified in (1), they actually can't constitute knowledge after all, and it was only because we were mistaken about their content that we ever considered them to be known. This objection, like those preceding it, also has the cost of divorcing motor action from knowledge. However, there are two additional problems that make this a weaker objection than the previous one. First, I think our intuitive judgement is that these beliefs are indeed known, even when we specify their (slight) inaccuracy and inaccessibility. Returning to the example of D.F., I think it's clear we want to say she knows object position, size, and orientation when she successfully acts on them, even though she has no reflective access to this information. Moreover, I don't think stipulating less noise than precludes successful action changes this judgement.

However, even if it turns out this intuition is not widely shared, we still have a compelling reason to resist accepting (C). Specifically, it opens the door for a form of skepticism, on which a wide swath of our putative knowledge would likely fail to be known. First, as noise is a product of the physical limitations of our perceptual system (Faisal et al. 2008; Dhawale et al. 2017), we might reasonably expect it for all demonstrative mental content that amounts to a belief of one's perceptions. Examples of this include belief that an object is that colour (as in Hellie 2005) or beliefs involved in perceptual judgements regarding an object's location (as in Fiehler et al. 2010; Odegaard et al. 2015). In these sorts of cases, belief content will be determined by noisy perceptual content, which cannot be known under factive knowledge. I will discuss this in much more detail in chapter five.

Moving beyond perceptual belief, we find an independent skeptical threat in our representation of numerical magnitude. Just as in the case of perception, although perhaps more surprisingly, magnitude representation is also variable and noisy. As I will discuss at length in chapter five (§4), it is widely accepted that our representations of numerical magnitudes, which guide numerical discrimination tasks and arithmetic calculation are inherently noisy (see Ansari and Vogel 2013), and they display the same trial-to-trial variability (Feigenson et al. 2004) we see in motor action and perceptual judgement. It even appears our representations of numerical magnitudes vary with gaze orientation (Brunamonti et al. 2012) and visual experience (Pasqualotto et al. 2014). All this suggests an argument similar to the one presented in this chapter, except this time it is not our perceptual knowledge of the spatial properties of objects under action that is threatened, but our *a priori* knowledge of numerical magnitude.

The fact that accepting (1)-(3) and (C) opens the door for this skepticism is itself a reason to look for a safer objection to my argument. Additionally, like those objections discussed previously, it still manages to divorce our concept of knowledge from action. As it appears to violate our intuitive judgements for knowledge and invites skepticism about a non-trivial proportion of our putative knowledge, while still accruing the same costs as arguing that the beliefs I'm interested in are uniquely beyond the interest of epistemology, this doesn't seem the most viable route.

To conclude, in this section I have looked at three ways one might respond to my argument, in terms of (2) and (C): One might simply reject (2), accept (C), or argue that the belief descriptions I take for (2) fall outside the domain of epistemology because they involve reflectively inaccessible content. This third objection appears the most viable, and comes with the consequences of both divorcing epistemology from motor action and distancing epistemology from empirically informed descriptions of mental content. In the next section, I will argue that there is a fundamental problem with preserving any theoretical constraint on knowledge at such a cost. In this manner, we might understand rejecting factivity as the best solution to the problem introduced in section two.

4. Retaining Factivity

Up to this point, I have explored the costs and problems associated with four potential responses to the argument introduced in the first section: First, one might simply reject assumption (1), which has the obvious cost of divorcing our account of knowledge from the empirically informed description of mental content. The primary problem with this move is that it seems we might employ conceptual arguments to do similar theoretical work. Next, one might reject assumption (2) outright. This comes with the costs of abandoning our ordinary practice of belief attribution and divorcing action from belief (and therefore knowledge). Moreover, this move has the problem of being unable to generalize beyond beliefs that guide action. After this, I looked at what I consider the strongest objection to my account, the "exclusionary response" that we shouldn't require a theory of knowledge to handle beliefs with inaccessible content. This too has the cost of divorcing knowledge from both action and the empirical description of mental content, with a potential problem of accounting for why these demonstrative beliefs are beyond the purview of the epistemology of knowledge. Finally, one might simply accept (C). Beyond the cost of divorcing action from knowledge, this response comes with the significant problem of opening the door for a form of skepticism, not just of perceptual knowledge but even *a priori* mathematical knowledge.

With all this in mind, this section will argue that rejecting factivity is the best response to the problem presented in section one. In order to do this, I will proceed in the following way: First, I will make clear exactly what I mean by "rejecting factivity," and from this discuss the cost associated with this move, which is more or less simply that it is revisionary. After this, I will contrast this cost with the two reoccurring costs from other responses, divorcing knowledge from (i) action and (ii) empirically informed descriptions of mental content. In both cases, I will argue that a very similar problem emerges if we accept either of these costs in order to preserve the factivity of knowledge: Accepting either cost on the basis of preserving factivity fundamentally misunderstands the theoretical work we intend to be done by a necessary condition of knowledge.

First, I want to discuss what I mean with "replacing factivity" or "non-factive knowledge." I do not have in mind that under a non-factive theory of knowledge, all justified (etc.) beliefs would suddenly be counted under knowledge, or anything like that. More generally, I do not have in mind that any beliefs we previously judged to be unknown would now count as knowledge under a non-factive conception, or that any beliefs judged to be known would now fail to qualify as knowledge.[45] It is important not to take non-factive knowledge to be the same sort of claim as, for example, factive justification, on which whether you commit to factive justification will determine which beliefs you take to be justified (for example, Sutton 2007). As I understand it, not only might non-factive knowledge capture our current intuitive judgements regarding knowledge, it would allow us to do exactly that without insulating epistemology from empirical descriptions of mental content.

As discussed in chapter one, the most attractive feature of factivity is this ability to capture and explain our intuitive judgements regarding knowledge. It provides the basis for a simple theoretical structure of knowledge and does a significant amount of theoretical work in explaining which beliefs are known and which are not. However, there doesn't appear to be any reason why factivity is the only condition on knowledge that could do this same work. Imagine we replace it with a pseudo-factive accuracy condition, on which the standard of accuracy is extremely high, as I do with the "truthlikeness" condition in the next chapter. Such a condition, which would yield a type of analysis we might call "factivoid," would certainly preserve everything known under factive knowledge, as anything that fulfils factivity will also fulfil pseudo-factivity. Moreover, if we require accuracy on par with that of the content that guides successful motor action, it doesn't seem any ordinary, coarse-grained beliefs would be at risk of now being mistakenly counted as knowledge. There is of course, as with any analysis of knowledge, the worry that we might imagine some contrived counterfactual in which the factivoid conception over-attributes knowledge, and I will have to let this worry stand for now.[46] Nevertheless, at least initially there doesn't seem to be any *prima facie* reason to think that something like a pseudo-factive accuracy condition couldn't do

---

[45] Of course, it is likely that this isn't entirely preventable. For any given non-factive analysis of knowledge, we might anticipate some contrived cases that that analysis is unable to handle. However, as this also holds for factive accounts of knowledge, this would hardly be attributable to replacing factivity.

[46] The only problem cases I can foresee will involve approximately true beliefs in the exact results of precise measurements (like the "$2\ell_P$" example from section three). However, it doesn't sound unreasonable to posit that owing to the approximating nature of measurement itself, it is never appropriate to form "exact" beliefs on the basis of measurement. Thus, we might say these beliefs are never justified and (hopefully) preëmpt any problems they might cause for my nascent non-factive account of knowledge. I discuss this more in the next chapter.

the same theoretical work as factivity. Again, I will discuss this matter at length in the next chapter.

Accordingly, the cost of rejecting factivity cannot be understood in terms of a loss of theoretical work, or a failure to capture our intuitive judgements. Rather, the cost is that of any nascent revisionist theory: It is presently unclear what exactly a factivoid theory of knowledge might look like, and determining this will require serious philosophical effort. Many, I'm sure, might be skeptical that such an effort would even be successful. Furthermore, were it successful, it would then dictate the revision of all those epistemological accounts that presuppose that knowledge is factive (which is to say, most of them). This is certainly a non-trivial cost. However, I think it would be a profound mistake to accept either of the rival costs, (i) or (ii) above, as a preventative measure. In both cases, the reason is rather simple: The presumptive purpose of necessary conditions on knowledge is to track either our intuitive judgements or perhaps some objective features of cognition.

Let's begin with the problem of accepting a divorce of action from knowledge. While this is also a cost of rejecting (2), accepting (C), here I will focus on the cost as it specifically holds for the exclusionary response, as I think this is the strongest route one might take. In any case, the terms of this "divorce" might be thought of as something like accepting (C). If the beliefs that guide motor action are to be excluded from our theories of knowledge on the grounds that their content is partially inaccessible, then of course those beliefs cannot be known. Furthermore, if, as I also argued in the previous section, the demonstrative beliefs that guide our motor action are the same ones that provide reason for motor action, there is a further divide between action and knowledge. Not only does it follow that the beliefs that guide successful motor action cannot be known, but the beliefs that provide reason for rational motor action also cannot be known. In this way we can understand the full extent of separation between knowledge and action: The beliefs that guide successful motor action and provide reason for rational motor action cannot be known.

The problem associated with accepting this cost is that this goes against both our general intuitions and intuitive judgements regarding the relationship between knowledge and action. These are beliefs that we intuitively judge to be known, and intuitively we want to say that our theories of knowledge should be able to accommodate these beliefs. This problem is most pronounced if we take something like the "Action-Knowledge Principle" as an epistemic norm, on which you should "treat the proposition that p as a reason for acting only if you know that p" (Hawthorne and Stanley 2008, 577). For the case of motor action, the results would be devastating. Every reason for motor action, being unknowable, would necessarily violate this principle. Of course, we need not commit to "Action-Knowledge" for this to be a problem. Whether we account for rational action in terms of knowledge, warrant (Gerken 2011), or its irreducible self (Levy 2013), the common thread is that, intuitively, the beliefs that guide and provide reason for acting (but especially the latter) are ones that are ordinarily known, or at least knowable.

Of course, one might respond that even if we need to sacrifice our intuitions and intuitive judgements with respect to knowledge and action, this is ultimately worth it because it allows us to retain the factivity condition. However, put in these terms, we can now see the fundamental problem with such a move. The very reason we posit necessary conditions for knowledge is to capture our intuitions and intuitive judgements. We expect these conditions to covary with our judgements, and in cases they don't, we should not exclude recalcitrant beliefs, but rather take them as a reason to revise those conditions. To suggest that we privilege a condition like factivity over our intuitions and intuitive judgements fundamentally misunderstands the relationship between the two.

I think the only clear sense in which this might be incorrect is if we don't take knowledge to be some extant concept, the nature of which is evidenced by intuition and intuitive judgement, but rather some kind, natural or otherwise, much closer to a scientific than folk concept. Kornblith's account knowledge as a natural kind, discussed in chapter one, presents a good example of this, as Kornblith explicitly rejects the idea that intuitive judgement has a major epistemological role to play (2002). However, if you go this route, it's unclear how you might possibly accommodate the cost of divorcing knowledge from empirical accounts of mental content. Rejecting empirical claims on the grounds of retaining factivity is obviously antithetical to such a naturalistic project. Moreover, even on a weaker version of epistemological naturalism, like Goldman's, it is still unimaginable that one might be willing to accept the cost of divorcing epistemological from the empirical description of mental content. In short, whether one commits to analytic or naturalistic assumptions about the study of knowledge, we can understand that, given the associated costs, it is fundamentally mistaken to retain the factivity of knowledge.

5. Folk Knowledge Attributions

The above argument has relied on the assumption that the analysis of knowledge should seek to preserve our ordinary knowledge-attributing practices. It is on the basis of our ordinary knowledge attributions that I conclude both that (i) we seem to judge that the false beliefs that guide motor action are still known and (ii) we face a skeptical problem if we don't consider these beliefs to be known. However, there is a concern here that I place too high a value on the preservation of these epistemic folk practices, a concern that is likely heightened by the very naturalism that my argument requires. After all, there is a body of empirical evidence that suggests that our ordinary practices of knowledge attribution can sometimes be quite troublesome, leading some like Gerken (2017) to propose an epistemological method on which these practices play a more constrained role. It is tempting to simply chalk this up to a fundamental disagreement in method. After all, it's not as if the primacy I'm granting ordinary knowledge attributions is especially unusual. However, I think simply resting my argument against factivity on the methodological assumption that we fully preserve folk attributions leaves the argument in a rather weak position. In the light of an independent empirical case against the full incorporation of folk attributions, the more natural conclusion would seem to be that the familiar methodological assumption is flawed, and that I've simply misidentified a novel argument for restricting the epistemological role of folk judgements as an argument against factivity.

The purpose of this section is to address this type of concern. In order to do so, I'll follow Gerken's account somewhat closely, first looking at the troublesome folk practices he identifies, and then the methodological revisions he proposes. In doing so, I intend to demonstrate that this sort of account does not pose any real problem for my argument against factivity. In short, our ordinary practice of attributing knowledge in the cases of the false beliefs that guide motor action are not problematic like other folk judgements might be. Therefore, even following Gerken's proposed restrictive method, the folk practice in that case is not easily discarded.

Gerken's methodological recommendations derive from three "shifty patterns of knowledge ascriptions" we might observe empirically[47]:

> The first one is an *alternatives effect*—roughly, the inclination to deny S knowledge that *p* in the face of a salient alternative, *q*. The second effect is a *contrast effect*—roughly, the idea that whether an alternative, *q*, to S's knowledge is "in contrast" partly determines our inclination to ascribe knowledge. The third effect is a *practical factor effect*—roughly, the effect of salient practical factors on our inclination to ascribe knowledge. (ibid, 20)

In order to illustrate each of these effects, I'll formulate a version of each specifically involving knowledge attributions associated with motor action.

> ***Base Example***: At a coffee shop, under ordinary circumstances, Sanna repeatedly reaches for and successfully grasps her cup of coffee. Does Sanna know where her cup is?
>
> ***Salient Alternatives Example***[48]: At a coffee shop, under ordinary circumstances, Sanna repeatedly reaches for and successfully grasps her cup of coffee. However, if a cunning prankster (1) secretly replaced Sanna's glasses with ones that that made it appear to her that the cup was actually much closer than it was and (2) without Sanna wising to the prank, expertly placed the cup in Sanna's hand whenever Sanna went to pick it up, Sanna wouldn't know where her cup is. As it happens, Sanna hasn't checked the veracity of her glasses, nor her surroundings for such a prankster. Does Sanna know where her cup is?
>
> ***Contrast Effect Example***[49]: At a coffee shop, under ordinary circumstances, Sanna repeatedly reaches for and successfully grasps her cup of coffee. There are other cups on the table that look

[47] See Shaffer & Knobe (2012), Nagel et. al (2013), Buckwalter (2014b), and Buckwalter & Schaffer (2015) for empirical evidence for the alternatives effect; Shaffer & Knobe (2012) and Gerken & Beebe (2016) for the contrast effect, and Pinillos (2012) and Sripada & Stanley (2012) for the practical factor effect.

[48] This is based on the format of Nagel's (2010) adaptation of Cohen's (2002) table example.

[49] This is at least inspired by Gerken's (2017) presentation of Schaffer and Knobe's jewellery thief example (2012).

indiscriminable from her cup. Does Sanna know where her cup is? Does Sanna know what cup is hers?

> ***Practical Factor Example***[50]: At a coffee shop, under ordinary circumstances, Sanna repeatedly reaches for and successfully grasps her cup of coffee. Rather unexpectedly, an evil scientist bursts through the doors, with a cup-grasping robot in tow. This robot is designed to grasp at coffee cups, guided by representations of a mug's location transmitted into it from an external source. The evil scientist announces that he has taken Sanna's entire extended family hostage. The only way they will be released safely is if the robot successfully grasps Sanna's coffee cup on the first try, guided by a real-time uplink of Sanna's dorsal-stream representation of the mug's location. Does Sanna know where her cup is?

Were these all paradigmatic examples of their respective effects, we might expect to not judge that Sanna has knowledge in light of the salient alternative, judge that Sanna doesn't know what cup is hers despite knowing where her cup is, and judge that Sanna doesn't know where her cup is when the safety of her entire extended family is at stake. Of course, I'm not going to claim that here I've captured the troublesome judgements sometimes associated with these effects. My intention here was simply to highlight their respective formats. Even at this point we might note that our attributive practice in the case of motor action is nothing like these three effects, which I'll discuss more shortly.

On the basis of the troublesome patterns of folk judgement displayed in these three effects, Gerken suggests that the role of folk judgements in epistemology needs to be constrained. He introduces multiple methodological principles on this basis, the most relevant of which is the "Principle of Principle:"

> Intuitive judgments provide significant evidence only if they can be theoretically integrated in a principled manner. Similarly, rejections of robust patterns of intuitive judgments must be principled. (2017, 60)

At this point, there are two distinctions to be made between our folk attributive practices associated with motor action and those that Gerken highlights. First, it is a rather simple matter to integrate these practices into epistemological theory. As I demonstrate in the next chapter, not only is the non-factive analysis theoretically viable, but it is even theoretically *advantageous*, most notably allowing us to unify the epistemological and scientific uses of "knowledge" under a single concept. Moreover, knowledge attributions associated with motor action are far more ubiquitous and robust than those in the alternatives, contrast, and practical factor effects. The fragility of folk judgements under these effects is underscored by my apparent inability to reproduce them in the "coffee shop" examples. We might easily contrast this with our attributions involving motor action. Under ordinary circumstances, we simply judge that the beliefs that guide successful motor action

[50] This isn't really based on anything.

are known. This is not some rare and elusive "effect" elicited by a narrow range of conditions, but a familiar fact of daily life.

In short, I think we can confidently conclude that the folk knowledge-attributing practices I seek to preserve in arguing against factivity simply do not display the hallmarks of folk practices we might want to revise. They are robust and ordinary, and the analysis of knowledge might accommodate them without too much trouble. Moreover, as we'll see in chapter five, we find equally robust and ordinary folk practices for both perceptual judgement and numerical magnitude, which also violate the principle of factivity. All this suggests that these are exactly the sorts of practices we want necessary conditions on knowledge to track. They are not anomalies we might easily discard.

6. Naturalistic and Folk Epistemologies

To close out this chapter, I want to briefly address an additional worry one might have about my methodology. The argument against factivity I've presented requires a blend of naturalistic and folk commitments. On the one hand, we need to accept a highly empirical description of mental content, but on the other hand we also need to accept our folk practices of forming judgements about this content. The tension between these two commitments is immediately obvious, and might lead someone to question why we should seek to marry naturalistim and folk practices in the first place. Perhaps the apparent problem for the factivity of knowledge is only a vestige of an unholy marriage between two disparate methods.

My response to this sort of concern is simply that I don't think that either naturalistic or folk elements are necessary to construct *an* argument against factivity. First, if we abandon folk practice entirely and commit to a singularly naturalistic account of knowledge, it is almost certain that this account will need to be non-factive, given that this is how "knowledge" is used is psycho-cognitive research (see chapter four, section three). Next, I think that our folk tendency to attribute knowledge to beliefs that are very close to true, guide successful motor action[51] but have minute, fine-grained error is independent of any empirical description of mental content. As I argued above (§2), I think that the theoretical work done by "Two Visuals Systems" and planning noise might also be done by conceptual argument, albeit perhaps not as convincingly. In short, I think that one might construct an argument against factivity either only on the basis of epistemic folk practice or only on the basis of a hard-line epistemological naturalism.

This then seems to raise the question of why I've opted for an argument that utilizes both naturalistic and folk components. The simple answer here is that my intention is to advance the strongest argument against factivity that I can devise. Moreover, as discussed in chapter one, I think both empirical findings and folk practices have key epistemological roles to play, and it is when we consider them together that the need for a non-factive analysis of knowledge comes best into focus.

[51] Or perceptual judgements and arithmetical computation, as we see in chapter five.

7. Conclusion

As it currently stands, the above argument against factivity has two key weak points. First, I have not satisfactorily established that factivity might be replaced with a weaker condition that might do the same theoretical work, or even discussed much of what this condition might actually look like. This is the primary task of the next chapter. Second, I have not satisfactorily established that there is a skeptical problem associated with retaining factivity. This is the task of chapter five.

To conclude, I will simply restate what I have argued here: According to a substantial body of evidence, the mental content that guides motor action slightly misrepresents the world. On the simplest read, this means that the beliefs that guide motor action are never true (i.e. anti-factive). However, as we still want these beliefs to count as knowledge, we must replace factivity with something weaker that can accommodate these beliefs. Various objections to this picture all either come at the cost of ignoring relevant empirical research and/or preventing the beliefs that guide and provide reason for action from ever being known. Accepting these costs on the basis of retaining a constraint on knowledge is a fundamental mistake, as we expect such constraints to track our intuitive judgements and (possibly) properties of cognition, not the other way around. Accordingly, we can understand the best option in this case is to accept that visuomotor noise precludes factive knowledge.

Chapter 4: The Factivoid Analysis of Knowledge

In the previous two chapters, I advanced an empirical argument against the factivity of knowledge. A large part of this argument rested on the conjecture that a non-factive analysis of knowledge might do the same epistemological work as a factive one. While I hinted at what such an analysis might look like, the details were far from complete. The purpose of this chapter is to fill in these details. Here I will look at three candidate analyses of knowledge, discussing them in order from most to least conservative. The purpose of these discussions is not to find a perfect analysis of knowledge, as such a result is beyond my present capacity to deliver. Rather, the intent of this chapter is to demonstrate that what I will call "factivoid" analyses of knowledge are *viable*. That is, an account of knowledge might shift from factive to factivoid without encountering any new problems with the under-attribution or over-attribution of knowledge.[52] Moreover, as non-factive knowledge allows us to accommodate both empirical data and ordinary attributions in cases like motor action, it is clearly preferable to a factive conception. Additionally, as will become clear in the final section, it also allows for more flexible theories of knowledge, which might accommodate both epistemological and scientific usage of "knowledge."

This chapter assumes a straightforward structure. After a brief, preliminary clarification on what I mean by "factivoid," I will simply discuss each candidate analysis in turn. Each section will be devoted to one such analysis, and structured in the same manner. First, I will introduce the analysis and give background on its origins. After this, I will first consider problems that that analysis might have with under-attributing knowledge, followed by a consideration of problems it might have with over-attribution. In each case, I intend to show that moving from factive to factivoid results in no new under or over -attribution of knowledge. Finally, in section four, I return to claim that factivity is needed to explain specific *linguistic* intuitions (introduced in section two of chapter one).

The first analysis I will discuss is justified truthlike belief. This is the most conservative analysis of the three, the primary purpose of which is to demonstrate that a non-factive constraint on knowledge might do all the same theoretical work as factivity. After this, I introduce a non-factive version of robust virtue epistemology. This analysis also serves to illustrate the viability of a non-factive replacement for factivity. Additionally, it also offers an explanation for the sort of non-factive constraint required. Finally, on an especially speculative basis, I use the framework of robust virtue epistemology to offer a broader account of knowledge, which is able to unify both epistemological and psycho-cognitive senses of "knowledge."

[52] To be clear, a theory T *under-attributes* knowledge for proposition p just in case (i) we judge that p is known and (ii) T entails that p is unknown. Conversely, T *over-attributes* knowledge for p just in case (i) we don't judge that p is known and (ii) T entails that p is known.

0. "Factivoid" vs. "Non-Factive"

Before detailing candidate factivoid analyses of knowledge, I first want to clarify what I mean by "factivoid" and distinguish it from "non-factive." As a non-factive state is simply a state that can link an agent to either truth or falsehood, a non-factive analysis of knowledge is simply an analysis on which knowledge that p does not entail p. In order for a non-factive analysis of knowledge to even seem plausible, we must place certain restrictions on which sorts of falsehoods might be known. The restriction I will use here is that, in addition to true propositions, only *truthlike* false propositions might also be known (see §1.1). Roughly, false beliefs might be known, but only if they are very close to the truth. A *factivoid* account of knowledge is one for which only true or truthlike propositions are known. So defined, a factivoid analysis of knowledge is a type of non-factive analysis of knowledge, seemingly the only type that even makes it past the smell test.

1. Justified Truthlike Belief

In this section, I will introduce the first of three candidate non-factive analyses of knowledge, justified truthlike belief (JTlB). As this is nothing more than a modification of justified true belief (JTB), it is the most conservative of the three non-factive accounts I will discuss in this chapter. Since it is based on JTB, it will of course retain all the same problems of JTB. Nevertheless, JTlB as a theory is instructive because it illustrates the viability of weakening factivity specifically and factivoid knowledge generally. We might substitute in a weaker condition for factivity in a way that doesn't introduce any new under or over -attribution of knowledge. The idea here is that if we hold the justification and belief conditions fixed, we can observe that a properly formulated non-factive account can do all the same work as a factive account, while also handling the cases introduced in the previous chapter.

In discussing the analysis of knowledge as justified truthlike belief, this section will proceed in the following way: First, I will introduce the truthlikeness condition on knowledge, and explain what it might mean for a belief to itself be truthlike. After this, I will consider possible problems with under and over -attribution. As all true beliefs will be truthlike, there isn't a risk of under-attributing knowledge. However, there is a worry that it might over-attribute knowledge in cases, like those noted in the last chapter, of approximately true false beliefs that are more reflectively accessible. I will argue at length that such cases don't actually present a problem for even this, the simplest non-factive account.

1.1 The Truthlikeness Condition

The most straightforward way we might adapt to non-factive knowledge is to replace the factivity constraint with a weaker necessary condition. My proposal for such a condition is simply a truthlikeness condition, which I will describe here.

The beliefs that guide motor action are problematic for standard epistemological theory because of their asymptotic relationship with the truth: They are often very close to truth. Learning processes might reduce noise, getting them even closer to the truth (see Dosher and Lu 2017). Under ordinary conditions we even judge them to be known. Nevertheless, given the coarse-grained, binary nature of the concept of truth, they still count as false. In order to capture this proximity to the truth displayed by some false beliefs and not others, we need a more fine-grained concept. The best option I see here is to simply go with a concept that corresponds roughly with the "distance" of a proposition from the truth. Thankfully, one such concept, *truthlikeness*, has already been described.[53]

First coined by Karl Popper (1963), "truthlikeness" is used interchangeably with verisimilitude as a gauge of the approximation of truth. Popper first introduces the concept as a means of evaluating the extent to which scientific theories correspond with the truth:

> I believe that we simply cannot do without something like this idea of a better or worse approximation to truth. For there is no doubt whatever that we can say, and often want to say, of a theory t_2 that it corresponds better to the facts, or that as far as we know it seems to correspond better to the facts, than another theory t_1. (ibid, 232)

Popper takes truthlikeness to be not just a property of scientific theories, but of individual statements as well, defined as "something like the difference of the truth-content and the falsity-content" of a statement (ibid, 393). One key feature highlighted by Popper is that truthlikeness is scalar and continuous. The truthlikeness/verisimilitude of statement a, $Vs(a)$, might assume any real number between -1, the truthlikeness of a contradiction, and $+1$, the truthlikeness of the truth (ibid, 391-7). That is, $-1 \leq Vs(a) \leq +1$. While we need not commit to the details of Popper's account[54] to utilize truthlikeness in this manner, this is the basic idea of truthlikeness I will use going forward, a scalar measure of verisimilitude.

Since truthlikeness is scalar, like justification, we might say that certain beliefs are more truthlike than others. Thus, for in order for the concept of truthlikeness to work as a condition of knowledge, we must set a specific threshold of truthlikeness necessary for knowledge. The result is the truthlikeness condition of knowledge.

> TRUTHLIKENESS CONDITION FOR KNOWLEDGE: S knows that p only if p is sufficiently truthlike.

In order for this to be a plausible replacement for factivity, the threshold for sufficient truthlikeness needs to be very high. Otherwise, any analysis of knowledge built on such a condition will encounter intractable problems with over-attributing. While we might simply leave it at that, I think there is an obvious standard for truthlikeness we might posit, the truthlikeness of successful cognitive function in

---

[53] More recently, Elgin has argued for the "epistemic acceptability" of false propositions that are "true enough" (2017, 9; see also 2004). While Elgin doesn't think that "true enough" propositions might be known, the concept of "true enough" is notably similar to that of truthlikeness.

[54] For a detailed description of this and other theories of truthlikeness, see Oddie (2016).

ordinary conditions. As we will see below (§1.3), the truthlikeness displayed by perceptual beliefs that guide successful motor action appears to be a viable standard for a necessary condition on knowledge. Finally, we might note that in replacing factivity with truthlikeness, we retain the intuition that there is a close relationship between knowledge and truth. In adopting a non-factive analysis of knowledge, we need not abandon the idea that there is a, as put by Buckwalter, "powerful and important relationship between knowledge and truth" (2014a, 407).

1.2 Under-attribution

Given that truthlikeness is a weaker condition than factivity, it doesn't appear JTlB will pick up any problems with under-attributing knowledge as the result of truthlikeness itself. If a belief is true, then it will be sufficiently truthlike. Therefore all beliefs known on JTB will also be known on JTlB.

1.3 Over-attribution

JTB is especially famous for problems with over-attributing knowledge in Gettier cases (1963), and it also appears to over-attribute in cases of probabilistic evidence (Hawthorne 2004; Friedman and Turri 2015). While JTlB will inherit these shortcomings, here I am only interested in *novel* over-attribution, cases in which the truthlikeness condition itself is responsible for JTlB over-attributing knowledge.

In the previous chapter, we observed that our intuitive judgements seem to differ for different cases of almost-true, fine-grained belief content. While we judge that ordinary perceptual beliefs are known, even if their content is partially determined by noisy, reflectively inaccessible motor representations, the same doesn't hold for approximately true beliefs that might emerge in the use of some high-precision measurement device, resulting in reflectively accessible but slightly inaccurate belief content. If this is correct, it appears to allow us to imagine cases in which justified reflectively accessible false belief A is more truthlike than justified, reflectively inaccessible false belief B, but A is unknown while B is known. If belief B is something like the ordinary perceptual beliefs that guide motor action, it seems belief A would then count as knowledge under JTlB. In this manner, we could construct novel over-attribution cases for the truthlikeness condition. However, as I will argue here, there is a problem if we try to actually find these beliefs. Given the requirements of constructing truthlike beliefs that we judge to be clearly unknown, these beliefs will never be sufficiently justified to qualify as knowledge.

In order to argue this point, this subsection will proceed in the following way: First, I will work through a couple examples until we get to beliefs with reflectively accessible content that are more truthlike than that of the perceptual beliefs that guide successful motor action, which we clearly judge to be unknown, taking care to note the parameters required for such beliefs. Most significantly, it only appears we might form such beliefs on the basis of measurement. After this, I will present two arguments why these sorts of beliefs will never count as knowledge. First, it seems we might describe them as inferential beliefs that rely on a false premise. Second, and more generally, it doesn't appear it is ever appropriate to form this type of belief on the basis of measurement, as they will never be justified. Accordingly, we can account for our judgements without having to posit some weird, likely multipartite condition that covaries with reflective access. We judge these beliefs to be unknown because one can never be justified in believing them. Thus, there is no over-attribution problem for the truthlikeness condition.

In order to understand what sort of cases might prove difficult for JTlB, let's begin with an example of a belief like those detailed in the previous chapter, which we judge to be known despite being only approximately true:

> Randy reaches to grasp the coffee mug in front of him. Accordingly, we attribute a perceptual belief to Randy regarding the mug's location, a belief that guides his motor action: Randy believes that the mug is there. The content of this belief is partially determined by a reflectively inaccessible motor representation of the mug's location. While this representation is more than sufficiently truthlike to successfully guide Randy's motor action, it is slightly noisy. Accordingly, it represents the mug as being at a location it is not. Nevertheless, our intuitive judgement is that when he initiates his movement, Randy knows that the mug is there.

If we wish to consider a belief that approximates truth in this way, but is reflectively accessible, the first thing we might try is simply imagining we make our motor representations reflectively accessible. Consider we continue Randy's story in this vein:

> . . . Unbeknownst to Randy, an evil scientist has spiked Randy's coffee with a drug that induces new neural pathways in the brain, providing the taker reflective access to all previously inaccessible dorsal-stream and low-level representations. This access is similar to memory retrieval. While not present in a constant stream like visual experience, if Randy reflects on this information he might access it. Upon reaching for his coffee mug a second time, he does just that. Now for Randy's belief that the mug is there, Randy can report the content of "there" in terms of three-dimensional egocentric coordinates. Again, planning noise results in this belief not being perfectly accurate. However, one side effect of this drug is that the accuracy of motor representations is drastically increased. This reflectively accessible content is much more accurate than the previous, inaccessible content, resulting in a more accurate belief.

This I think is an odd case, and I can't say that I have a strong intuition one way or the other. I'm tempted to say that Randy still has knowledge, but it's not clear whether this is an intuitive judgement anymore or simply the product of my theoretical commitments. However, as it seems likely that the murkiness emerges from the strange mental content, we might remove this to construct a case in which we judge highly truthlike but false beliefs to be unknown. As the content in question will need to be both reflectively accessible and very accurate without being perfectly accurate, the only obvious source of such beliefs seems to be quantitative measurement. If the content isn't quantitative, it's unclear how it might fulfil the accuracy requirements. Similarly, without relying on measurement, it's also unclear how we might achieve sufficient accuracy while maintaining reflective access and some ostensible source of justification.[55] Accordingly, we might imagine something like the following:

> Jan, a benevolent scientist, invents an extraordinary device that can determine and display the **exact** location of any object, down to the Planck length (ℓ_P), in any conceivable reference frame. After getting Randy's neural pathways and perceptual capacities back to normal, she uses this device to measure the location of the coffee mug in the same egocentric coordinates as the inaccessible motor representations that guide Randy's motor actions on the mug. However, immediately after doing so, a freak, imperceptible earthquake shifts the mug $1\ell_P$ radially towards Randy. Thus, when Jan looks at the display screen and forms the belief that the mug is exactly at (r, θ, φ), with the content of this belief being determined by her reflectively accessible representation per her measurement, her belief content is far more accurate than Randy's. However, it is still not true, only truthlike.

Despite Jan's belief being far more accurate than Randy's, it seems we intuitively judge that she doesn't know that the mug is at (r, θ, φ), whereas Randy does know where the mug is when he acts on it. Given that Jan's belief appears truthlike, and at least initially appears to be justified, this poses a challenge for the truthlikeness condition and JTlB. It seems this might be a case of novel over-attribution on the part of JTlB. Here I will describe two responses we might give to this sort of problem: We might (1) maintain that knowledge cannot be inferred from a falsehood, or (2) note that it seems inappropriate to ever form an exact belief on the basis of measurement.

One might at first be tempted to abandon truthlikeness, or any unitary condition that captures accuracy, and instead maintain that there are different standards of knowledge depending on whether certain content is reflectively accessible. However, this would certainly result in an undesirably two-tiered replacement condition for factivity and strikes me as unsatisfyingly *ad hoc*. Furthermore, I think we have good reason to suspect that reflective accessibility isn't the relevant feature here. Besides the matter of reflective accessibility, there are other distinctions we might draw between Jan and Randy's respective beliefs.

One critical difference between Jan's and Randy's respective beliefs about the location of the mug is the source of their respective belief content. Randy's belief content is entirely perceptual, in this case the direct result of his visual capacities. However, Jan's belief content, as it is determined by a measurement, is not entirely perceptual. Rather, it results from an inference she makes from her

[55] The only other option seems to be we posit exotic mental content as a counterfactual. However, as illustrated in the second vignette (above), this seems to come at the expense of the clarity of our intuitive judgements.

perception of that displayed measurement, along with a facilitating implicit assumption:

(P1) The measuring device displays the position of the mug as being (r, θ, φ). [*perception*]

(P2) If in its current configuration the measuring device displays the position of the mug as being (r, θ, φ), then the mug is exactly at position (r, θ, φ). [*implicit assumption*]

(C) The mug is exactly at position (r, θ, φ).

In this manner, we can understand that to the extent to which it is determined by the measuring device, Jan's belief that the mug is there will be inferential. However, the problem with this inference is that it is based on a false premise. When Jan forms her belief, P2 is false. It is not the case that the machine displaying position (r, θ, φ) entails that the mug is at position (r, θ, φ). If we posit "no false premises" as a necessary condition for knowledge, one might appeal to this to explain why she lacks knowledge, without needing to posit separate epistemological standards for knowledge depending on reflective access.

While I think some might find this response attractive, this account is not without its limitations. First, it does still require we place an additional constraint on knowledge itself. Although it doesn't seem as arbitrary as positing separate constraints on knowledge for accessible and inaccessible content, I still think this leaves something to be desired. Next, while I won't discuss it here, I should at least note that a number of authors have recently argued that we can indeed infer knowledge from false belief (Warfield 2005; Fitelston 2010; Luzzi 2010; for a recent response see Schnee 2015). Finally, this still does seem to set up multiple standards for knowledge, a factive standard for inferential knowledge and a non-factive standard for perceptual knowledge. Again, while this divide looks less arbitrary than positing a non-factive standard for knowledge involving noisy, reflectively inaccessible content, and a factive standard for (tentatively) everything else, it still strikes me as a unsatisfying. I think we can do better.

There is a more intractable problem with cases like Jan's belief (quantitative, fine-grained, extremely but not perfectly accurate, reflectively accessible content formed on the basis of a measurement) than the potential for unreliable judgements or being based on a false premise. In order to understand this problem, we first need to note two necessary features of these cases so that they might compare to noisy perceptual content: They must be based on measurement, and their content must be exact. They must be based on measurement because there appears no alternative means by which we might rationally acquire beliefs with fine-grained error comparable to that caused by noise in human perception. Of course one might base such beliefs on nothing at all, but these won't pose a challenge for JTIB because they won't be justified. Furthermore, the content of these beliefs must be exact because if it is not, we judge them to be known and there is no longer any complication for a non-factive account. Consider if instead of exact representation of location, the content of Jan's belief is less determinate. If her reflectively accessible representation is that the mug is somewhere close to (r, θ, φ), and this is what determines the content of her belief that the mug is at (r, θ, φ), we simply judge this belief to be known.

However, the problem that emerges for cases with both these properties is that it never seems appropriate to form an exact quantitative belief on the basis of measurement.[56] At an intuitive level, it seems we judge a speaker is doing something wrong if they report an exact belief on the basis of a measurement. Imagine someone uses digital callipers, precise to the nearest one tenth of a millimetre, to measure the diameter of a cylinder. Upon doing so, they report the diameter to be exactly 5.17 cm. When pressed, they don't concede that this report wasn't meant to be exact, and perhaps the diameter isn't exactly 5.17 cm, but just some value close to it. They insist on the basis of their measurement that the diameter of the cylinder is exactly 5.17 cm. Even before we fill in the details, there seems to be something wrong with this belief. Furthermore, even if we specify that this person used the callipers competently, and that the object happens to be exactly 5.17 cm, it seems their belief is profoundly mistaken. In fact, it seems that even granting that the belief is true, we don't judge it to be known.[57] At the very least, this will hold for any actual case involving exact belief and measurement. This I think is what our intuitive judgements are responding to when we judge that Jan's belief is unknown. It is simply in the nature measurement that it does not provide an adequate basis for beliefs with exact quantitative content. Or, put another way, these beliefs are not justified.

This observation that it seems we are never justified in forming exact beliefs on the basis of measurement corresponds well with standard metrological[58] theory and what has been called the "epistemology of measurement" (Mari 2003)[59]. It is widely recognized that any actual measurement will be subject to inevitable uncertainty from multiple sources, including environmental factors and limitations on the instrumentation (ISO 1993). Additionally, uncertainty is now generally considered to be fundamental to the nature of measurement itself, the product of the modelling required of measurement:

> Against the tradition, which has considered measurement able to produce pure data on physical systems, the unavoidable role played by the modeling activity in measurement is increasingly acknowledged, particularly with respect to the evaluation of measurement uncertainty. (Giordani and Mari 2012, 2144).

More generally, as put by Mari (2003, 24) this metrological stance on uncertainty reflects the "need to make any empirical process of information acquisition

[56] To be clear, here I'm talking about exact beliefs in the quantities being measured, not the measurements itself. Of course on the basis of measurement we might believe that the measuring device displays exactly (r, θ, φ), etc., but this is something different.

[57] This might be explained by noting that these sorts of beliefs aren't sensitive to the fact of whether the quantity in question is exactly X. Consider the closest possible worlds in which diameter was actually exactly 5.17000000000000001 cm. The callipers would still display "5.17," still resulting in beliefs that the diameter was exactly 5.17 cm.

[58] Metrology is the academic study of measurement.

[59] Although I think "philosophy of measurement" is a more appropriate description of the endeavour.

approximate." Exactly true values are viewed as simply unattainable via measurement[60] (ibid, 24). This serves to corroborate our initial suspicion that exact beliefs formed on the basis of measurement will never be justified. Accordingly, they will not be considered known on JTlB, which means that there is no threat of over-attribution.

I would suspect that our initial tendency to not question that these beliefs are justified arises from the fact that ordinarily measurement is uncontroversially a source of justification. In ordinary practice, given competent measuring procedure, the beliefs we form on the basis of measurement tend to be justified. However, ordinarily we don't form beliefs that the world is exactly as the measurement reports, and upon consideration it appears to be a violation of competent measuring practice to do so.

The most obvious way one might object to this picture is to maintain that measurement is not the only plausible route to over-attribution for JTlB. The best way I can imagine one doing so is with a case like the following:

> Martha, a physics student, reads in her text book, "The fine structure constant α is $\frac{1}{137}$." On this basis, she comes to believe that the value of the fine structure constant is exactly 1/137. However, the fine structure constant is not actually exactly 1/137, but only approximately 1/137 (1/137 - α $\approx$ 0.000002). Somehow in the editing of the textbook, the "approximately" was erroneously removed from the text.

Here we might note that while false, Martha's belief is highly truthlike, and it seems she's justified in believing it. However, we certainly don't judge that Martha knows that the fine structure constant is 1/137. In this manner, it appears that in cases like these JTlB over-attributes knowledge. While this might be correct, I don't think that such cases are *novel* instances of over-attribution. It is a matter of luck that the typo in the textbook resulted in a truthlike sentence, and therefore it is a matter of luck that Martha's belief is truthlike. As JTB famously struggles with this sort of knowledge-undermining luck, I don't think there's anything novel about JTlB over-attributing here.

In conclusion, there doesn't appear to be any novel risk of JTlB over-attributing knowledge. The only potentially justified beliefs that are accurate enough to count as known under the truthlikeness condition, but still false, will involve beliefs in the exact values of measurement. However, upon closer inspection it doesn't appear that such a belief might ever be justified. This both explains why we judge them to be unknown and eliminates the initial problem that they appeared to present for JTlB.

[60] There is a very interesting parallel here between this view of measurement and the picture of vision-for-action detailed in the previous two chapters. Mari even hints that "sense data obtained by direct observation" will also have the uncertainty endemic to measurement (2003, 24).

1.4 Conclusion

In this section, I hope to have demonstrated the basic viability of a factivoid analysis of knowledge. Taking JTB as a starting point, we might replace the truth condition with a truthlikeness condition without encountering any novel under-attribution or over-attribution of knowledge, provided that our standard for truthlikeness is set at the accuracy of ordinary perception. However, one limitation of simply positing the truthlikeness condition is that it is unable to account for why perceptual accuracy should determine standard of truthlikeness. While doing so enables a viable factivoid analysis of knowledge, it remains a bit of a mystery why this actually works. In the next section, I will introduce a more sophisticated factivoid account, which not is not only viable but is able to explain why setting the standard of truthlikeness to that of ordinary perception proves successful.

2. Factivoid Robust Virtue Epistemology

In this section, I will discuss the second of three candidate non-factive analyses of knowledge, a factivoid version of robust virtue epistemology (RVE). I intend to show that not only is the structure of RVE is especially amenable to the shift away from factivity, but that it might explain why the standard of truthlikeness for knowledge should be determined by the accuracy of ordinary perception. As with the previous section, this one will be structured in the following way: First, I will introduce a factivoid version of RVE, after which I will consider possible problems each might have with the novel under and over -attribution of knowledge.

2.1 Factivoid RVE

The purpose of this subsection is to introduce a factivoid version of robust virtue epistemology. In doing so, I will first briefly outline the standard view, followed by a transition from standard to factivoid RVE.

Coined by Pritchard (2010), "robust virtue epistemology" refers to the claim that "knowledge is type-identical with a particular species of cognitive achievement" (Carter, Jarvis, & Rubin 2015, 1603). As achievement here is simply success attributable to ability, we might understand RVE as the claim that knowledge is cognitive success attributable to cognitive ability. The distinguishing feature of *robust* virtue epistemology, as opposed to virtue epistemology generally, is that "it attempts to ***exclusively*** analyse knowledge in terms of" a cognitive success that is attributable to a cognitive ability (Pritchard 2010, 25; emphasis mine). Of course, this alone is more a template for an analysis of knowledge than it is an analysis itself, as such a claim is at this point ambiguous in multiple respects. One such ambiguity involves what it means for a success to be "attributable" to an ability, but this won't be discussed here.[61] This is simply because factivoid RVE doesn't need to favour any particular take on attribution. Next, we might inquire about the specific nature of a cognitive ability. On the

[61] For an account of the various proposed versions for this attribution relation, see Greco 2012, §3.

standard account, "a cognitive ability is a knowledge-conducive belief-forming process" (Pritchard 2012, 261), and I won't diverge from this here. Rather, here I will disambiguate one key question: What constitutes a cognitive success? By answering this in a non-standard way, specifically a non-factive way, we might construct a viable factivoid RVE. First, however, we should begin with the standard epistemological account of cognitive success.

On the standard picture, true belief is cognitive success and false belief is cognitive failure. If a belief is true, it constitutes a cognitive success; if a belief is false, it constitutes a cognitive failure. To avoid confusion, we should note that the sense in which "cognitive" is used here is not the sense used elsewhere in this project. It is not the sense in which we talk about *cognitive* systems or *cognitive* neuroscience, which encompass far more than the formation of belief. Rather, this is a sense of "cognitive" is restricted to the domain of belief, and might be interchanged with "doxastic" (Sosa 2009, 10), "epistemic," (Turri et al. 2017) or "intellectual" (ibid). While in section three we will revisit this restriction, here I will only talk of cognitive success insofar as it is doxastic success.

At this point, we might question why we should consider true belief to be a success and false belief a failure. The standard answer to this question is that the *aim* of belief is truth:

> The aim of belief is said to be truth. When you sincerely pose a question to yourself, for example, you want a correct answer. When you reach an answer to your question through adopting a certain belief, the aim of your belief is the truth of the matter. (Sosa, 2009, 6)

This view that the aim of belief is truth has been challenged along the lines that it doesn't actually explain why "correct beliefs are true beliefs" (Owens 2003, 283; see also Whiting 2014; for a response see Steglich-Petersen 2009, 2017). However, this idea that belief aims at truth faces another problem in that it doesn't seem to reflect the actual properties of belief-forming processes. This problem might first be observed as an intuitive matter. When we consider beliefs like those that guide motor action, it doesn't appear that they are aiming at truth at all. Upon close inspection, it seems more appropriate to say that these beliefs aim at a certain level of truthlikeness. This is especially clear when we note that certain belief-forming processes, like those for action guidance, lack the *capacity* to hit truth. While in ordinary circumstances this process produces truthlike beliefs, the inevitability of visuomotor noise makes it extremely unlikely that any such belief will ever be strictly true. Accordingly, it seems mistaken to think of these beliefs as aiming at the truth. Any appearance of aiming at truth here is only due to the proximity of the actual aim, truthlikeness, to truth. Given that at least some of the time belief aims not at truth but only truthlikeness, we can understand that it is mistaken to take only true belief to constitute cognitive success.

This idea, that often beliefs aim not at truth but a standard of truthlikeness, and thus might be successful without reaching truth, gains further support if we turn our attention to the *output requirements* of belief-forming processes. Recall that in chapter three we briefly encountered the concept of the output requirement of a cognitive process. As Goodale and Milner (1992) noted, the ventral (vision-for-conscious-perception) and dorsal (vision-for-action) streams are required to output different types of representations. Roughly, the ventral stream is required to output more qualitative representations (colour, texture, etc.) while the dorsal stream is required to output more quantitative spatial information (location, size, etc.). Additionally, while the output of conscious perception is required to be conducive to long-term storage in memory, the opposite holds for vision-for-action (see especially Milner and Goodale 1995, §2.4). All this raises an obvious question, "Required for what?" While Milner and Goodale don't state explicitly, the answer seems to be something along these lines: For the ventral stream, these are the required outputs for successful conscious perception. For the dorsal stream, these are the required outputs for successful motor action. Thus we might say, as a general matter, the output requirements of a cognitive process are those required outputs for cognitive success.

Here, the type of output representations required for success isn't directly relevant. However, I think we might extend the concept to the output requirement to the accuracy of beliefs as well: There is a certain degree of accuracy, a standard of truthlikeness, required of the output of belief-forming processes in order for them to successfully constitute conscious perception, guide action, etc. Therefore, we might understand the success of any given cognitive process as determined by the output requirements of that process. Put simply, cognitive success is determined by output requirements.

This idea, that what it means for a belief to be successful is determined by the output requirements of the processes on which that belief forms, provides us with a foundation for a factivoid RVE. Upon noting that often truth is not an output requirement of belief-forming processes, we might understand cognitive success in the non-factive sense required. Cognitive success is not truth, but the standard of truthlikeness determined by the output requirements of a given belief-forming process. This leaves us with a theory of knowledge on which knowledge is still cognitive achievement, with the following conceptions of success and ability: Cognitive abilities are still taken to be knowledge-conducive belief-forming processes. However, for certain knowledge-conducive belief-forming processes, such as vision-for-action, cognitive success is not truth but rather a sufficient degree of truthlikeness. In this way, we might describe a factivoid RVE.

Before continuing to the potential for novel misattribution for factivoid RVE, there are two things I would like to note. First, we might observe that on this account *cognitive ability determines cognitive success*. That is, what it means for a belief to be successful is not fixed independently of the process by which it forms, but is determined by the output requirements of that process. As cognitive abilities are simply those belief-forming processes that are knowledge-conducive, we might understand that what it means for a belief to be successful is determined by the ability by which it forms. While I am hesitant to make any definitive conclusions, this at least opens the possibility of cognitive success being constituted by different standards of truthlikeness for different cognitive abilities. For example, it might be that for certain cognitive abilities, perhaps those processes associated with inferential knowledge, the output requirement is still truth. This of course in an empirical matter, and I won't go into it here.

The second observation we might make is that, unlike an analysis like JTlB, which simply posits a truthlikeness condition, factivoid RVE provides an explanation for why the standard of truthlikeness corresponds with the truthlikeness of ordinary perception. This is the minimum accuracy output requirement we might observe for knowledge-conducive belief-forming processes. Achieving that standard of truthlikeness because of the operation of a knowledge-conducive belief-forming process is simply achieving cognitive success because of cognitive ability, which is simply achieving knowledge

2.2 Under-attribution

One of the major challenges posed to RVE is that it under-attributes in cases of knowledge from testimony (Lackey 2007, 2009). To the extent that this is a problem for standard RVE, it will also be a problem for factivoid RVE. However, as with Gettier problems for JTB, here we must set these concerns aside in the interest of isolating novel under-attribution.

As with the move to a truthlikeness condition presented in the previous section, there doesn't appear to be any risk of novel under-attribution for factivoid RVE. The reason for this is similar in both cases. Just as all true propositions will be truthlike, all cognitive success on standard RVE (true belief) will count as cognitive success on factivoid RVE (true or truthlike belief, determined by the specific output requirements of individual belief-forming cognitive processes). Since cognitive success on factivoid RVE is a weaker concept than cognitive success on standard RVE, all cognitively successful beliefs on the latter will also be successful on the former. Therefore, there is no risk for novel under-attribution.

2.3 Over-attribution

RVE appears to struggle with over-attributing knowledge in certain cases involving knowledge-undermining luck (see Pritchard 2012, 15-17). However, again, in the interest of identifying over-attribution associated with the move to a factivoid account, I'll set this aside.

Despite my best efforts, unlike JTlB, there doesn't appear to be any obvious candidate for novel over-attribution on factivoid RVE. Cases involving truthlike measurement don't pose any problem, as there is no cognitive ability by which we might form exact beliefs on the basis of measurement. As discussed in the last section, it appears to violate competent measuring practice to believe physical quantities are exactly what one measures. Therefore, the belief-forming process by which one does so clearly does not constitute a cognitive ability. Also like JTlB, it may be that factivoid RVE struggles with certain beliefs that are truthlike as a matter of luck. However, if this were indeed the case, it would of course not be novel over-attribution.[62]

[62] There might be a worry that that because the candidate factivoid analyses I discuss all have been challenged by cases involving epistemic luck, that perhaps there is something inherent to factivoid analyses that they cannot handle luck. However, this worry might be easily addressed by noting that we might easily construct a factivoid anti-luck condition, factivoid safety: If S knows that p then S's truthlike belief that p could not have easily been not-truthlike (see Prichard 2012, §2). This allows the

2.4 Conclusion

Not only does factivoid RVE again demonstrate the viability of the factivoid analysis of knowledge, but it also provides valuable insight into why the standard of truthlikeness for knowledge appears to be set at the level of ordinary perceptual accuracy. Of course, factivoid RVE will inherit all the shortcomings of standard RVE, but this doesn't pose any problem for the factivoid analysis of knowledge itself.

3. Radical Factivoid RVE

In the previous section, I presented a factivoid version of robust virtue epistemology. On this account, knowledge is still cognitive success because of cognitive ability, but cognitive success is understood in such a way that certain false beliefs might still constitute a cognitive success. The key idea there was that it is the output requirements of belief-forming processes that determine what it means for a belief to be successful. In this section, I will take the brakes off this analysis and utilize the RVE framework to a more radical effect. The proposed account of knowledge will still be cognitive achievement, but I will no longer restrict the domain of cognitive success to belief formation. Bringing the structure of RVE in line with a more standard use of "cognitive," cognitive achievement need not involve belief formation at all. This is the least conservative of the three accounts discussed in this chapter, and it brings with it the consequences we might expect from this. However, dropping the doxastic restriction from RVE also affords significant advantages. Most notably, it allows for the unification of the epistemological and psycho-cognitive usage of "knowledge."

In the familiar structure, this section will first introduce radical, factivoid RVE, followed by a discussion of the motivation behind such a move. After this, I will consider potential concerns of either under-attributing or over-attributing knowledge. Finally, I will look at the necessity of dropping factivity and the ability to do the same for belief. Again, in reading this section, it is important to keep in mind that this proposal is especially speculative, and in places will push at certain assumptions of this project.

3.1 Domain Unrestricted Cognitive Achievement

As with the initial version of factivoid RVE discussed in section two, let's once again choose as a starting point the idea that knowledge is cognitive achievement, cognitive success because of cognitive ability. This time, however, rather than restricting "cognitive" to be synonymous with "doxastic," let's instead recognize that many of what we would otherwise call "cognitive processes" or "cognitive systems" do not reduce to belief formation. On radical factivoid RVE, cognitive success might be constituted by any cognitive process, not just belief-forming

---

factivoid analysis of knowledge to account for cases like the fine structure constant example (§1.3). Her belief was formed on the basis of a typo, and a typo could have just as easily resulted in a belief that wasn't truthlike at all.

ones. Again, the question of what constitutes success is an empirical one, determined by the output requirements of individual systems.

Immediately, this proposal seems unnecessarily broad in at least two respects. First, on this analysis belief is no longer a necessary condition for knowledge. For the moment I will set this concern aside. Later in this section I will argue that this might be considered advantage of the theory, although of course it brings with it certain costs. Second is the even broader concern that I'm no longer just talking about propositional knowledge. Given that many cognitive processes might be successful without generating propositional knowledge, the initial risk of over-attribution seems extremely high. This is of course correct. However, the fact that such a proposal includes more than just propositional knowledge doesn't mean it is a bad account of knowledge. Rather, I think it offers the opportunity to unify multiple senses of "knowledge" under a single concept. The core of the proposal is this: Knowledge is a unitary concept we might analyse as cognitive achievement, and different senses of "knowledge" correspond with specific restrictions on what can constitute a cognitive success.

In order to illustrate this, here I will identify three robust senses of knowledge: propositional knowledge (knowledge-that), practical knowledge (knowledge-how), and the more empirically oriented concept we might see in cognitive science and psychology, "psycho-cognitive" knowledge. In each case, we might understand the specific sense of knowledge as a restriction on the domain of cognitive success. Depending on what we take to constitute a cognitive success, we will get different senses of knowledge.

PROPOSITIONAL KNOWLEDGE: As a starting point, we might say that propositional knowledge is knowledge for which the domain of cognitive success is limited to propositional content. For propositional knowledge, cognitive success is successful formation of propositional content, and cognitive failure is failure in the formation of propositional content. Of course, this framework is ambiguous both in terms of what constitutes propositional content and what constitutes the successful formation of propositional content. The first ambiguity we might ignore here, as it is no more ambiguous than for any other epistemological account of propositional knowledge, so it's not clear that the disambiguation of "propositional content" is of immediate relevance. As long as we understand propositional content as in some vague sense representing the world to be a certain way, this will suffice for the sketch of this account. The second ambiguity we might fill in. Just as for the more constrained factivoid RVE proposed in section two, cognitive success is determined by the output requirements of the relevant cognitive abilities. If a specific cognitive process requires a high degree of truthlikeness, then cognitive success for propositional content produced by that process is that high degree of truthlikeness, etc. Again, what constitutes cognitive success is in this sense an empirical matter, and there is at least a possibility that this will vary for different cognitive processes.

While I will reserve the following sections for discussions of the potential for this conception to under and over -attribute propositional knowledge, there are two sub-categories of propositional knowledge I wish to introduce here, which might preëmpt certain worries about over-attribution. These I will call "doxastic knowledge" and "theoretical knowledge[63]." At the most basic level, ***doxastic knowledge*** is simply knowledge that is associated with a belief attribution. Given that belief is a propositional attitude, we can understand this to be a sub-set of propositional knowledge. Accordingly, if for some attribution of knowledge that p we can identify a corresponding belief that p, that propositional knowledge is doxastic knowledge. If we wish to consider belief to be a real mental type, then we might take this a step further and say that doxastic knowledge is a knowledge for which the domain of cognitive success is limited to belief formation. Next, ***theoretical knowledge*** is knowledge with reflectively accessible content. Given the robust empirical distinction between accessible and inaccessible mental content, we might formulate theoretical knowledge as knowledge for which the domain of cognitive success is limited reflectively accessible propositional content. Theoretical knowledge is the knowledge we might report on, use in our reasoning, and so forth.[64]

Before continuing, there are a few observations I'd like to make regarding the relationship between propositional, doxastic, and theoretical knowledge. First, while there is certainly a good deal of overlap, doxastic and theoretical knowledge are distinct categories. As illustrated by the mental content that guides motor action, knowledge might be associated with a belief attribution without retaining accessibility. Thus, not all doxastic knowledge is theoretical knowledge. Next, at least initially, it seems we want to say that all theoretical knowledge is doxastic. This keeps us in line with standard epistemological practice, which is desirable since there is no obvious reason to break with it. Finally, while we might observe that all doxastic and theoretical knowledge will be propositional, as will become clear shortly, there may be some propositional knowledge that is neither doxastic nor theoretical.

PRACTICAL KNOWLEDGE: First, let me clarify that here I will use "practical knowledge[65]" to denote what is often called "knowledge-how." Examples of practical knowledge range from simple motor ("S knows how to grasp a coffee mug.") and perceptual skills ("S knows how to discriminate between red and green.") to more complex skills, which might range from the very broad ("S knows how to play football.") to the very specific ("S knows how to run a post route against single coverage."). The reason that I am not simply using "knowledge-how" is that sometimes we use "knows how" to attribute propositional knowledge, unrelated to knowledge of some skill or ability. For example, we might say, "Jukka knows how the 2010 MSU-Notre Dame game ended" to denote clearly propositional knowledge: Jukka knows that the 2010 MSU-Notre Dame game ended with MSU faking a field goal to score a game-winning touchdown." To filter out this sort of propositional knowledge, it's easier to just go with "practical knowledge."

[63] After a similar notion from Glick (2011, 412)

[64] As discussed in the previous chapter, in some cases the contents of beliefs might be determined by both accessible and inaccessible content. Accordingly, it is best to think of theoretical knowledge admitting of degrees: Some instances of knowledge are more theoretical than others, being more reportable, integrable into reasoning, and so forth.

[65] This term is also use this way by Glick (2011).

While some, most famously Stanley and Williamson (2001), have argued that practical knowledge is simply a form of propositional knowledge, there is compelling empirical evidence that propositional and practical knowledge are distinct. First noted by Cohen and Squire (1980), neuropsychological evidence suggests that practical knowledge dissociates from propositional knowledge in the following way: Amnesiac patients can acquire a new skill (on the basis of which we would attribute practical knowledge) and retain that knowledge even after they forget all propositional knowledge associated with the acquisition of that skill. The example they use involves the skill of reading words printed as if viewed in a mirror ("mirror reading"). They found that not only might individuals with amnesia might be taught how to mirror read, but also that they retain this knowledge even after losing all propositional knowledge associated with it.

These results led Cohen and Squire to posit a distinction between the representation of "declarative" and "procedural" information (ibid, 209), a distinction that serves as the foundation for the contemporary empirical argument in epistemology for the independence of propositional and practical knowledge (Wallis 2008; Adams 2009; Brown 2013). At its core, this distinction is between two independent memory systems, a procedural memory system for "mediating the acquisition of cognitive skills" and a declarative memory system for "mediating learning of new facts or other data-based knowledge" (Cohen et al. 1985, 69). When we gain new skills, the corresponding representations are stored in procedural memory, fundamentally distinct from the propositional knowledge stored in declarative memory:

> Procedural representations are those that guide cognitive operations and overt actions by specifying what is to do under which circumstances (where circumstances are given through declarative representations). Procedural representations can be explicated as condition-action rules, with the condition describing the circumstances to which the procedure applies, and the action component describing what is to be done (i.e., the cognitive operation or the physical action to be carried out). Although communicating procedural representations (e.g., when instructing people) requires translating them into a declarative format, such as a verbal rule or a graph, they are not themselves declarative: A core assumption of theories distinguishing declarative and procedural representations is that cognition and action is controlled by procedural representations, and declarative representations such as instructions must first be interpreted by procedures that take them as input and generate (cognitive) actions as output. (Oberauer 2010, 279)

This distinction between declarative and procedural representations allows us to identify, at least preliminarily, the domain restriction for practical knowledge: Practical knowledge is cognitive achievement for which the domain of cognitive success is limited to procedural content. When the procedural memory system produces a representation that meets its output requirements, it achieves practical knowledge.[66]

PSYCHO-COGNITIVE KNOWLEDGE: Even when framed in terms of domain restrictions for cognitive success, the categories of propositional and practical knowledge are both epistemologically familiar. However, in addition to these categories, there is a third sense of knowledge that appears to be just as robust. This is the sense of knowledge we might find associated with empirical research into human cognition. Here I will rely primarily on examples from cognitive neuroscience, especially cognitive neuropsychology, but I don't mean to claim this is a uniquely cognitive-neuroscientific usage. This usage corresponds with what we might call "psycho-cognitive" knowledge. While here I won't endeavour to make any definitive conclusions on the boundaries of this concept, upon initial inspection it appears to correspond with the general analysis of knowledge I'm presenting here, domain-unrestricted cognitive achievement.

As a disclaimer, I should first note that the term "knowledge" is not used as consistently as I may make it seem below. Some researchers use it liberally, while others not at all. It is not uncommon to see a paper that outlines a claim about "knowledge," referencing a paper that instead uses a word like "representation" instead of "knowledge" in establish that claim (eg. see Gade et al. 2004, 174; referencing Squire 2004). Moreover, there is significant variability between research literatures with respect to the frequency of the usage of "knowledge," clustering in specific patterns for reasons that aren't entirely transparent. We might compare the proliferation of "procedural knowledge" (discussed below) and "prior knowledge[67]" (Faisal et al. 2008, 298) with the conspicuous lack of "knowledge" in the Milner and Goodale program (chapter 2, section 3). While this might suggest a need for additional research, I don't think it detracts from the observations about the psycho-cognitive usage of "knowledge" I discuss below.

When we first examine the way in which the word "knowledge" is used in psycho-cognitive research, we find that it often corresponds with familiar theoretical knowledge. Knowledge corresponds with propositional content, facts about the world that the knower might access upon reflection. Belief reports evidence what an agent does and does not know. For example, the fact that a lesion patient "knew that a kangaroo is Australian and an elephant is not, despite his inability to report how many legs each has" is taken to evidence that there are separate systems for knowledge of perceptual properties and knowledge of non-perceptual properties (Rapp 2001, 6). Similarly, "knowledge" is used to refer to a known neuropsychological finding, such as "knowledge of the neural bases of bilingualism" (ibid, 322). In both cases, knowledge is used in the familiar sense of theoretical knowledge.

---

[66] I say that this is only a preliminary proposal because it is likely there are multiple cognitive systems that produce practical knowledge, with the perceptual learning associate with discrimination skills being distinct from procedural memory/content (see Squire 2004).

[67] Used in the context of neural strategies to minimize the impact of noise, "prior knowledge" refers to the brain's representation of expect levels and distribution of noise across perceptual channels (Faisal et al. 2008, 298; see also Van Beers et al. 1999; Körding & Wolpert 2004).

The interesting features of psycho-cognitive knowledge arise from the fact that "knowledge" is also used in a way that is markedly non-theoretical. First, psycho-cognitive knowledge need not be reflectively accessible, as in "knowledge without awareness" or "unconscious knowledge" (ibid, 89). In short:

> Contemporary psychology speaks of unconscious knowledge to refer to cases in which subjects display available knowledge to which they lack conscious access. While this is not controversy-free in psychology, a significant part of the psychological community attributes to this claim a scientific status. (Augusto 2011, 116)

Unconscious knowledge is attributed when "subjects exhibit behaviours that indicate that they possess knowledge but seem both unaware of that possession and unable to verbalize it" (ibid, 117). Perhaps the best examples of this come from unconscious perception, as discussed in chapter 3. In the case that S unconsciously perceives, say, the orientation of a grating, S can know the orientation of that grating without having access. Given that we wouldn't on the basis of S's unconscious perception attribute a belief to S, we might also understand that psycho-cognitive knowledge can be non-doxastic.

More generally, cognitive knowledge seem to include all cognitive success in the representation of information, regardless of the level of processing or degree of accessibility:

> In psycho-cognitive terms, knowledge can be defined as information or data about the environment (roughly: sensory input) that can be acquired, stored, and retrieved by living organisms with a more or less complex nervous system, with a view to securing their wellbeing. Cognition can be defined as the actual process of acquisition, storage, and retrieval of knowledge. (ibid, 117)

While such a definition serves to illustrate the breadth of psycho-cognitive knowledge, it of course is remarkably crude. Beyond obvious epistemological problems, one of which I discuss in §3.2, this still doesn't capture the full scope of the concept. Already in the examples above, we can see psycho-cognitive knowledge includes information that is neither about an agent's environment nor in any way related to her wellbeing, such as the lesion patient's knowledge that kangaroos are Australian. Beyond this, we might note that the representations associated with the procedural memory system also fall under psycho-cognitive knowledge as "procedural knowledge." The following are examples from, in order, sports psychology, cognitive psychology, and cognitive neuroscience (emphases mine):

> It is widely believed that highly practiced, overlearned performances are *automated*—meaning they are controlled in real time by ***procedural knowledge*** that requires little attention, operates largely outside of working memory, and is substantially closed to introspection. (Beilock & Carr 2001, 702)

> The ability to acquire declarative and ***procedural knowledge*** has been mapped to different neural substrates. (Gade et al. 2014, 174)

> Altogether, perceptual and motor skill learning paradigms indicate that a high degree of plasticity has been preserved in the adult human cortex and that a limited repertoire of neuronal mechanisms of plasticity mediate memory functions throughout the mammalian cortex. This plasticity may underlie the acquisition and the long-term retention of ***procedural knowledge***. (Karni 1996, 47)

In short, psycho-cognitive knowledge includes practical knowledge ("procedural knowledge") as well as both theoretical and non-theoretical, doxastic and non-doxastic, propositional knowledge. Accordingly, it appears that this empirical concept of knowledge is simply knowledge as unrestricted cognitive achievement. Analysing knowledge along these lines allows us to not only to account for both the psycho-cognitive and epistemological uses of "knowledge," but also to unify them under a single concept that explains the relation between them. While in the form presented this is simply the sketch of a theory, I think it offers a promising route for further development. In the following two subsections, I'll explore some of the challenges the theory faces at this stage, and in doing so begin to refine it.

3.2 Under-attribution

As noted above (§2.2), RVE struggles with under-attribution in the case of knowledge from testimony, famously noted by Lackey (2007, 2009). In the interest of only discussing the potential for *novel* under-attribution, this concern will be set aside.

At a first glance, it doesn't appear that radical factivoid RVE runs much risk of under-attributing knowledge in a way that standard RVE does not. After all, the former is a much broader conception of knowledge, which counts under it far more than the latter. This instinct, that the risk for under-attribution seems low when moving to a much more attribution-happy theory of knowledge, seems largely correct. However, there is in interesting sort of under-attribution this account appears to be guilty of, which I want to discuss here. Put briefly, there seems to be psycho-cognitive knowledge that does not constitute cognitive achievement.

Beyond the other issues with defining psycho-cognitive knowledge as information about an environment (discussed in §3.1), we might note that this definition leaves no room for a distinction analogous to that of true belief vs. knowledge. So defined, psycho-cognitive knowledge includes all environmental information, even that which is truthlike as a matter of epistemic luck. There are of course many ways in which one might gain information as a matter of luck, and we might even note that Gettier's original "Case I" (1963, 122) is an example of luck in the acquisition of environmental information.

While it is tempting to simply chalk this space for luck up to the non-analytic nature of such definitions, in this case I think this would be a mistake. In the case of non-environmental or non-beneficial knowledge, we might easily observe that psycho-cognitive usage does not actually ignore these categories. However, this does not appear to be the case for accidentally truthlike representation. Not only have I been unable to find anything to indicate that psycho-cognitive knowledge ever excludes information acquired as a matter of luck, but we might also observe that the way in which the concept is employed is often necessarily insensitive to sources of information. Let's consider an example of how "knowledge" is used in the neuropsychological studies of the patient known as "AC". While AC could report certain properties of objects, like whether an animal is Australian he struggled to report other properties, like which animals have legs (Sartori and Job 1988). Summarized one way:

> What happened to AC is that he lost all information about the visual properties of objects, while still retaining knowledge about their non-perceptual properties, and also knowledge about perceptual properties which are not visual. What does this tell us about how knowledge of objects is represented mentally? . . . [There] is a system of knowledge about what objects look like which is quite separate from other stores of knowledge about other kinds of properties of objects. (Coltheart 2001, 6)

What I want to highlight here is that even if we assume that AC's beliefs about the non-visual properties of objects are all truthlike as a matter of luck, it is clear these count as knowledge under this psycho-cognitive usage. Conversely, if we stipulate a conception of psycho-cognitive knowledge with an explicit anti-luck condition, it undesirably narrows the scope of this claim. The researchers here want to draw a conclusion about how visual information is processed, regardless of why that information happens to be accurate.

In short, when psycho-cognitive knowledge is conflated with information in the nervous system, independent of whether that information is truthlike as a matter of luck, we should take it at face value. This usage of "knowledge" makes no distinctions for how information is acquired, the psycho-cognitive equivalent to conflating knowledge with true belief. Therefore, we might understand an analysis of psycho-cognitive knowledge as cognitive achievement as under-attributing psycho-cognitive knowledge, leaving out representations that are luckily truthlike, which do not constitute cognitive achievements.

There are two ways we might respond to this. First, we might go with a revisionary response and accept this under-attribution as a positive refinement on the psycho-cognitive usage of "knowledge." This would allow "information" to continue corresponding with cognitive success, with "knowledge" being reserved only for instances of cognitive achievement. We might also note that this isn't that disruptive a revision, given that in the psycho-cognitive realm there is very little cognitive success that isn't also cognitive achievement. In contrast with epistemology, the focus on actual and/or ordinary cases means that psycho-cognitive knowledge isn't usually applied in cases involving knowledge-undermining luck. Nevertheless, I think it is best to avoid philosophical revisions of scientific usage if it might be helped. The better option appears to the non-revisionary route.

The second response we might provide to this apparent under-attribution is simply to say that as a concept psycho-cognitive knowledge is not type-identical with an epistemologically interesting concept of knowledge. Rather that domain-unrestricted cognitive achievement, psycho-cognitive knowledge is domain-unrestricted cognitive *success*. While ordinarily tokens of psycho-cognitive knowledge will be cognitive achievements, we might note that this will not always be the case. This allows us to unify epistemologically interesting knowledge and psycho-cognitive knowledge under a single concept, while still respecting the autonomy of the latter.

In conclusion, the potential for radical factivoid RVE to under-attribute psycho-cognitive knowledge involving lucky truthlikeness results in a critical refinement to the analysis: Psycho-cognitive knowledge is best understood in terms of cognitive success, not cognitive achievement.

3.3 Over-attribution

The question of over-attribution for radical factivoid RVE is a complex one, and I will not fully address it here. Initially there appears to be a significant potential for over-attribution arising from low-level, non-doxastic, unconscious representations (e.g. the retinotopic representations discussed in chapter two (§2). However, upon closer inspection this isn't as obvious as it might seem. Not only might we question whether these intermediary representations actually constitute cognitive successes, but it also doesn't appear that these cases are considered psycho-cognitive knowledge. For example, in one study involving consciously unresolvable but unconsciously perceived orientation gratings, researches remark that the participants "had no knowledge" of the unresolvable stimuli (He and MacLeod 2001, 475). All this is to say that fear of over-attributing knowledge in the case of inaccessible content that we don't associate with belief attributions or skill execution might be premature. Nevertheless, to be safe, we should assume that there will be at least some over-attribution in cases like unconscious perception, so that we might see why we shouldn't consider this type of over-attribution to be a problem.

Let's begin with an example of the sort of cases in which we might anticipate radical factivoid RVE to over-attribute: In a study investigating the representation of orientation in the visual cortex, Haynes and Rees observed that simply on the basis of fMRI of subjects, "we could then successfully predict which one of two oriented stimuli a participant was viewing, even when masking rendered that stimulus invisible" (2004, 686). This indicates that the orientations of the stimuli were represented in the visual cortex, despite being masked from conscious perception. Assuming that this constitutes an achievement of the relevant cognitive system, this would seem to be an over-attribution by radical factivoid RVE. There is no relevant motor action, belief report, skill execution, etc., and it unclear that such representation might be capable of guiding any behaviour that we associate with knowledge. Accordingly, we don't judge that the subject knows what the orientations of the stimuli are. More generally, we might observe that it doesn't appear we ever attribute knowledge on the basis of fMRI alone in passive viewing studies.

My thought here is that if the analysis misattributes knowledge in this way, we might reasonably accept this as a cost of an analysis that can accommodate psycho-cognitive knowledge. After all, it is difficult to imagine such a project succeeding without this variety of misattribution. Empirical science highlights both mental states and behaviours that are invisible in ordinary life. We shouldn't expect our intuitive-judgement-forming capacities, which are tuned to more conspicuous behaviours and mental states, to be all that well equipped at handling cases so outside ordinary conditions. Accordingly, it seems simply accepting the limitations of our intuitive judgements in these cases is not all that unreasonable, and seems in keeping with the especially naturalistic flavour of this sort of account.

3.4 Conclusion

To conclude my discussion of radical factivoid RVE, I want to observe that this unification of epistemological and psycho-cognitive knowledge is only achievable under a factivoid framework. Given the argument presented in previous chapters, it shouldn't be surprising that psycho-cognitive knowledge is non-factive. Many of the cognitive systems described by psycho-cognitive research might be successful without generating true representations of the world. Furthermore, we might observe this directly in explicitly non-factive uses of "knowledge" and "know" in psycho-cognitive research, on which the knowledge might not be perfectly accurate (emphases mine):

> ***Knowledge*** of our hands' positions in space is vital for everyday tasks, yet in the absence of visual feedback we systematically misjudge their position and size . . . The aim of Experiment 1 was to identify how accurately and precisely participants ***know*** the orientation of their left and right index fingers in the frontoparallel plane in the absence of visual information. (Fraser & Harris 2016, 3565 & 3568)
>
> To advance understanding of mental health literacy on two prevalent neurocognitive disorders, this study aimed to develop and validate the Brain Injury and Schizophrenia Awareness Scale (BISAS) and compare accuracy of ***knowledge*** of the functional effects of schizophrenia and TBI. (McKendry et al. 2014, 225)
>
> People tend to be overconfident, that is, they exaggerate the extent to which what they ***know*** is correct. (Fischhoff et al. 1977, 552)

This underscores the importance of a factivoid framework for any account that seeks to analyse psycho-cognitive knowledge. Regardless of whether such a project is pursued along the lines of RVE, a theory on which knowledge is factive simply cannot capture the psycho-cognitive usage of "knowledge." This is a key advantage of factivoid analysis of knowledge, especially within a naturalistic context.

There is still a good deal of work that needs to be done before radical factivoid RVE might be considered a complete theory of knowledge. Not only is a more precise description of its boundaries needed, as well as a more sophisticated account of psycho-cognitive knowledge, but we might also observe that at least some of the non-factive uses of "know" above seem suspiciously close to protagonist projection. However, as mentioned at the outset, my intention with this preliminary attempt at radical factivoid RVE was not to produce a complete, challenge-free analysis of knowledge, but rather to illustrate how the factivoid RVE framework opens the door for a promising naturalistic analysis of knowledge with features we haven't seen before.

4. Factivity and Linguistic Intuition

Until now, I have reduced the question of the viability of factivoid knowledge to the question of whether a factivoid account of knowledge might successfully account for our intuitive judgements regarding knowledge, without misattributing in any novel ways. This is based on the assumption that the theoretical work of factivity is limited to its ability to account for our epistemic intuitions. While this is broadly true, there is an additional explanatory role that factivity plays. As discussed in chapter one, factivity is widely cited as an explanation for specific linguistic intuitions of oddness or contradictoriness for utterances like "S knows that p, and not p."

Factivity easily explains why these utterances seem contradictory. If knowing that p entails p, then such utterances are simply self-contradicting. However, if knowledge only entails the truthlikeness of p, this type of explanation no longer works. This, then, might be cited as an objection to my broader argument: Here we have something factivity does that a non-factive condition can't do. My response to such an objection is simple: It is not necessary that knowledge *is* factive in order to explain these linguistic intuitions, but only that factivity has a general intuitive appeal, that knowledge *seems* factive. While factivity certainly explains the fact that these utterances sound contradictory, it is far too hasty to think it might only do so if knowledge actually is factive. In order to illustrate this, let's first consider another utterance that might sound odd or contradictory:

> "Väinö lives in Finland, but he doesn't live in Scandinavia."

Although it is not guaranteed, I suspect that many will find this utterance to seem odd or contradictory in the same fashion as, "S knows that p, and not p." However, in this case we might note that the explanation for this intuition is not, and moreover *cannot*, be that living in Finland entails living in Scandinavia. This is because, while often misidentified as such, Finland is not part of Scandinavia. In this case, it is obvious that this is explained by the fact that one thinks Finland is in Scandinavia, that it seems like Finland is in Scandinavia. Moreover, even if we substitute "a Nordic country" for Scandinavia, it would be clearly mistaken to infer from our intuition of contradictoriness that Finland is in fact a Nordic country. Rather, the appropriate inference would be that it intuitively seems that Finland is a Nordic country.

In case one might be concerned about the dissimilarity between the utterances in question and the Finland example, we find something very similar with utterances like the following:

"S knows that p, but might be wrong."

As with the previous utterances, this too seems contradictory. As put by Lewis:

> If you claim that S knows that P, and yet you grant that S cannot eliminate a certain possibility in which not-P, it certainly seems as if you have granted that S does not after all know that P. To speak of fallible knowledge, of knowledge despite uneliminated possibilities of error, just *sounds* contradictory. (1996, 549)

While knowledge being infallible would certainly explain this intuition, we of course don't need to posit that knowledge is infallible in order to explain it. Instead, we only need to note that knowledge *seems* infallible, that there is a general intuitive sense that knowledge is an infallible sort of thing. Crucially, we might reject infallibilism even while noting that utterances like "S knows that p, but might be wrong" sound contradictory.

I think we might extend this same sort of explanation to "S knows that p, and not p." These utterances sound contradictory not because knowledge is factive but simply because knowledge seems like a factive sort of thing. This, however, can be explained by knowledge being factivoid. If knowledge is factivoid, then knowledge appears factive. It is only close investigation that reveals that it isn't. Moreover, as discussed in chapter one, general intuitions, like the intuition that knowledge is factive, might be overridden by both patterns of specific intuitive judgements and other theoretical factors (e.g. avoiding skepticism). As the case against factivity has both, we might understanding how we can explain the seeming contradictoriness of "S knows that p, and not p" without undermining this argument.

Chapter Five: Generalization

The argument against the factivity of knowledge presented in chapter three is, of course, a highly specialized one. It focuses specifically on knowledge associated with the visual guidance of motor action, which admittedly represents a very small subset of our total knowledge. Because of this, the purpose of this chapter is to make efforts directed towards the generalization of this argument. As we will find, the question of exactly how much my argument against factivity generalizes is a complex one, and here I won't be able to answer it conclusively. However, we can say quite confidently that there is nothing anomalous about motor action. Empirical findings highlight that many of our knowledge attributions associated with perceptual judgement (§2) and proprioception (§3) also violate the principle of factivity. Most significantly, this evens seems to hold for much of our *a priori* knowledge of numerical magnitude (§4). These observations highlight the extent of the skeptical problem associated with retaining factivity. Not only does this strengthen the general case for a factivoid analysis of knowledge, but the specific argument advanced in chapter three, which relies in part on the claim that there is a major skeptical threat associated with denying knowledge in the case of the beliefs that guide motor action. While I'll be unable to provide a complete account, and some key questions will remain unanswered, in this chapter the extent of the skepticism invited by factivity begins to come into focus.

This chapter will assume a simple structure. After an introductory section outlining my "formula" for constructing arguments against factivity, I'll (loosely) apply this formula to visual perception, proprioception, and magnitude representation. The last of these cases is by far the most interesting, and I'll accordingly spend the most time discussing the representation of magnitude. Finally, I'll say a bit on what appears to be the most significant question governing just how far the argument against factivity generalizes, whether we might make a similar argument for beliefs with more accessible content.

1. How to argue against the factivity of knowledge in four easy steps!

Before presenting any parallel arguments against the factivity of knowledge, I'd like to provide a generalized sort of "recipe" for the type of argument that I've made and will continue to make. Note that the primary value of this recipe is clarificatory, and I'm not trying to suggest that this is some robust philosophical method. Presenting a general formula for an argument in this way is merely a tool that helps highlight key components of that argument. Accordingly, in this chapter I'll only follow the recipe loosely. Moreover, we might note that while the argument presented in chapter three also contains all the key components I'll discuss here, its was of course structured differently. Finally, I should also note that I don't mean to suggest that this is the only way one might go about arguing against factivity, although it's presently unclear to me how one might go about it in a different way.

The recipe I present in this section involves four steps, identifying in turn (1) appropriate mental content, (2) behaviour guided by that content, (3) associated belief attributions, and (4) associated knowledge attributions.

Step 1: Identify appropriate mental content

This first step is of course empirical, and for this project I've been relying primarily on findings from cognitive neuroscience. Starting with empirical research, the key is to identify the potential for fine-grained misrepresentation, mental content that might misrepresent the world, but only slightly. The idea here is that the misrepresentation needs to be so slight that it is not observable in ordinary settings, but only in laboratory environments. I've found success in cases that either involves some sort of cognitive bias or cognitive noise (or both). Evidence that suggests that this misrepresentation causes variation in behaviour is crucial, and thus behavioural evidence for noise and/or bias is especially valuable. Evidence from neuroimaging, neuropsychology, and neurophysiology can also be helpful, but only insofar as it can be linked to ordinary behaviour. We should be especially on the lookout for fine-grained misrepresentation that is unavoidable and/or does not preclude cognitive success, as this will be helpful in step 4.

Step 2: Identify Behaviour

Once promising mental content has been identified, we need to identify the behaviour that is guided by this mental content. This too relies on empirical findings, and is most easily accomplished if the previous step draws from behavioural studies. After all, if misrepresentation such as noise or bias is posited in order to explain a certain behaviour, that of course is compelling evidence that said behaviour is guided by noisy and/or biased representations. After specific behaviour guided by the promising mental content is identified, we can move on to step 3.

Step 3: Identify Belief Attributions

Here we identify what beliefs we might attribute on the basis of the behaviour identified in step 2. Of course, if the behaviour is too fine, such as slight delays in response time often observed in priming paradigms[68], it may be that there aren't any associated belief attributions at all. If this is the case, we should then return to step 2 to find new behaviour (or step 1 if there isn't any). Once associated belief attributions are identified, we should check the truth-values of the beliefs. Ideally, it should be uncontroversial that they are false. However, as we see especially in §4.2, this can be a quite complicated matter and might require a good deal of work to establish. At any rate, once it can be established that the beliefs are in fact false, we can move on to step 4.

[68] See, for example, the studies utilizing priming paradigms discussed in chapter two (§1).

Step 4: Identify Knowledge Attributions

The final step is to check to see if we attribute knowledge on the basis of the false beliefs identified in step 3. Of course, the best possible result here is that we simply judge these false beliefs to be known, even after becoming aware of their slightly misrepresentative content. However, even if this isn't the case, we might also have some strong theoretical motivation to insist that these false beliefs are counted under knowledge. As I have argued in chapter three and will continue to argue below, denying that these beliefs are known invites a skeptical problem. If no such reason can be found, then we should return to step 3 and try again with different beliefs. However, if there is a compelling case to be made that we should consider these false beliefs to be known, congratulations! You've successfully argued against factive knowledge!

Of course, the simplistic "recipe" presentation downplays the challenges faced by this sort of argument. Each step requires a serious amount of argumentation and might be subjected to any number of objections. Nevertheless, by outlining these steps, which I've found myself working through in the process of building this project, I hope to provide a bit of clarification on what seem to be the fundamental components of this type of argument. Moreover, where presented in this way, we can clearly observe the "multi-level" structure of my argument, first discussed in chapter one.

2. Visual Perception Revisited

The purpose of this section is to (slightly) expand the set of knowledge that we might observe causes a problem for factivity. It seeks to address the matter of knowledge of spatial properties associated not with motor action, but perceptual (especially visual) judgement. In doing so, I'll give an overview of the knowledge attributions I have in mind, after which I'll present empirical evidence that the corresponding mental content is both noisy and biased. In short, this mental content that guides perceptual judgements poses the same problem for factivity as the mental content that guide motor action.

The argument against factivity presented in chapter three hinged on our readiness to attribute knowledge on the basis of success in ordinary motor action. If under ordinary conditions, guided by vision, S successfully reaches for and grasps an object, we will ordinarily grant that she knew where that object was, how large it was, and how it was oriented. In short, up to this point, the following is an exhaustive list of the knowledge attributions that I've argued challenge factivity:

> LOCATION: On the basis of her perception that object O is there, S knows that O is there. (S knows where O is.)
>
> SIZE: On the basis of her perception that object O is that large/small, S knows that O is that large/small. (S knows what the size of O is.)

> ORIENTATION: On the basis of her perception that object O is oriented thus and so, S knows that O is oriented thus and so. (S knows what the orientation of O is).

If you recall from chapter three (§3), I've argued that during motor action the content of these beliefs is determined by both ventral (vision-for-conscious-perception) and dorsal (vision-for-action) -stream representations. The main benefit of this move is that it still allows for the same beliefs to explain both motor action and perceptual reports. As the dorsal-stream representations are noisy, this means that during motor action these beliefs will be false, although highly truthlike, and the epistemological leg of the argument takes over from there. However, we might note that motor action is not necessary for us to attribute knowledge of this type, as we often will do so on the basis of S's belief reports, or perhaps even her gaze orientation or fixation. In these cases belief content is not determined by dorsal-stream representations, and thus there is no visuomotor-noise-based problem for factivity. Nevertheless, as I will argue in this section, empirical evidence also suggests that ventral-stream representations are both noisy and biased. Accordingly, it appears that whether we make the above types of knowledge attributions on the basis of successful motor action, perceptual judgement, or gaze, this knowledge can only be handled by a factivoid account.

To begin with, we might note that the sources of noise are not unique to visuomotor processing. As mentioned in section three of chapter two, noise is the product of the basic properties of the physical processes that serve as the foundation for perception and cognition (see Faisal et al. 2008). Far from a quirk of motor action, noise is ubiquitous throughout the nervous system:

> Noise permeates every level of the nervous system, from the perception of sensory signals to the generation of motor responses, and poses a fundamental problem for information processing. (ibid, 292)

Moreover, it is worth noting that there is also evidence that "the very different neural processes of [conscious] perception and action are limited by the same sources of noise" (Osborne et al. 2005, 412), and specifically for "shared noise sources in the [conscious] perception and action pathways arising from a common sensory estimate" (Mukherjee et al. 2015, 8515). On the basis of these observations alone, we might anticipate that ventral-stream representations of the spatial properties of objects are no more accurate than their dorsal-stream counterparts. This expectation matches the observation of noise-correlated inaccuracy in perceptual judgements, summarized as follows (emphasis mine):

> Perceptual decision making is one of the most important tasks performed by the brain . . . The brain must therefore have evolved a sophisticated and rapid mechanism for making decisions based on input from the external world. Physical stimuli have certain, well-defined properties when they arrive at the sensory organs, but they are then transduced and transmitted along pathways that have multiple synaptic way-stations where information can be delayed, distorted or influenced by descending, efferent activity. The accumulation of these effects is referred to as 'internal noise' . . . ***Percepts,***

> ***or internal stimulus representations, are therefore imperfect representations of the physical stimulus, and physically identical stimuli can elicit variable percepts.*** The decision-making mechanism must act on the perceptual evidence available to achieve the organism's goals . . . ***As the relative contribution of the internally-generated noise to the percept increases, decision-making becomes increasingly prone to errors.*** (Amitay et al. 2013, 1)

There are a few key observations here I'd like to stress. The first is that the mental content that guides perceptual judgements is very much susceptible to noise, as is evidenced by the correlation between noise levels and error in perceptual judgement (ibid; Arazi et al. 2017). Moreover, this noise is ubiquitous in the mental content that guides perceptual judgements (Neri 2010; Dinstein 2015; Ratcliff et al. 2018). As with the visuomotor noise discussed in chapter three, this means that this mental content will always at least slightly misrepresent the world, and we might provide a similar reply to the objection that noise should be understood in terms of increased indeterminacy: If we assume that this content is so indeterminate that it is no longer misrepresentative, it can no longer explain observed variation in behaviour. In short, it is clear that the mental content that guides perceptual judgement is inaccurate in the same sort of way as the mental content that guides motor action.

Beyond the general inaccuracy caused by noise, there is a more specific inaccuracy we might observe for ventral-stream representations of object location. Put simply, the representations of object location used in perceptual judgement are systematically biased towards gaze direction:

> People are biased toward localizing visual targets where they are looking. For static objects this bias is too weak to introduce noticeable errors, but for moving targets the temporal uncertainty can be large enough for the bias to no longer be negligible. (Brenner et al. 2006, 822)

Dubbed "foveal[69] bias" by Sheth and Shimojo (2001, 331)[70], the evidence for this bias comes from behavioural studies involving direct reports of perceived object location. As the bias is far more pronounced for objects in motion than it is for static objects, it's most easily observed in perceptual reports of objects in motion. For example, one such study by Brenner et al. (2008) used a visual array of five concentric circles, for which participants fixated their gaze at a small black dot in the centre of the array. The target was a green dot that jumped around the array, changing locations four times a second, and participants were asked to report on the location of the dot (which ring it was on) at a specific time, indicated either by a visual (flashing light) or auditory (tone) cue (ibid). In simply comparing actual target location to reported target location, this study "found a clear bias towards selecting target positions near fixation" (ibid, 8). Similar results have been observed in other numerous other studies, involving objects in motion (Kerzel 2002; Brenner et al.

[69] The fovea is the center of vision on the retina.

[70] Note that Sheth and Shimojo were not the first to observe this effect, which was explicitly noted by Rauk and Luuk in 1980, with similar observations dating to as early as 1955 (Leibowitz et al. 1955). For a brief account of this older literature, see van der Heijden et al. (1999).

2006) as well as the remembered location of static objects (Sheth & Shimojo 2001; Fiehler et al. 2010; Bocianski et al. 2010; Fortenbaugh and Robertson 2011). In short, there is strong evidence for the foveal bias (for an overview see Odegaard et al. 2015).

The misrepresentation attributable to this foveal bias is certainly not the only sense in which accuracy for the conscious perception of moving objects appears to be significantly limited.[71] However, as already alluded to, this bias is of particular interest because it not caused by motion:

> In this case the origin of the bias is not to be found in the statistics of the scenes that we encounter in daily life. Its origin lies in the way in which visual information is processed within the eye and brain, with most neuronal resources being devoted to a small area on the retina, and eye movements directing this part of the retina (the fovea) towards selected parts of the scene. (Brenner et al. 2008, 2)

This means that the bias is shared by representations of the location of static and moving objects alike. Finally, and more critically for our purposes, under most circumstances this is bias is weak to the point of unnoticeability:

> Since we do not expect a strong bias towards small retinal eccentricities[72], we can only expect such a bias to become apparent when there is considerable uncertainty about targets' positions. (Brenner et al. 2008, 2)

This combination of ubiquitous and scarcely noticeable misrepresentation lends itself to a similar result to that found with noise. It also helps to demonstrate how error in otherwise reflectively accessible mental content might still not easily be accessed upon reflection. To close out this section, I'll show how both cases allows us to construct an argument roughly following the recipe outlined in the previous section, thereby broadening the case against factivity and for the factivoid analysis of knowledge.

Let's begin with the sorts of beliefs we might attribute to S on the basis of her perceptual judgements. It should be uncontroversial that, in ordinary conditions, if S makes a perceptual judgement about the location/size/orientation of an object, we readily grant that S believes that the object is there/ that size/ that orientation. Next, we might note that the content of these beliefs is determined by noisy representations of location/size/orientation, which means that these beliefs will be false. Additionally, representations of location will also be biased towards the centre of gaze, resulting in an additional layer of misrepresentation. Given that perceptual judgements both covary with noise and directly display foveal bias, we might preëmpt any objection that belief content isn't actually determined by noisy/biased representations. Thus, we might reasonably conclude here that all our

[71] In one especially interesting example, there is evidence that, for judgements about the location of a moving object, "perceptual localization is influenced by motion signals collected over 80 ms after a query is triggered" (Eagleman et al. 2007).

[72] Retinal eccentricity is the angle of an image on the retina from the centre of gaze (fovea). Thus, "a bias towards localising targets near where one is looking is equivalent to a bias towards small eccentricities" (Brenner et al. 2008, 2).

beliefs about these spatial properties of objects will be false. However, ordinarily these beliefs are of course still highly truthlike. After all, we generally assume them to be true, and it is only through close observation in an experimental setting that their (slight) falsity might be observed. As is the case when the content of these beliefs is determined by noisy motor representations, they are false in just the right way to be problematic for the factivity of knowledge.

At this point, we might return to the cases of knowledge attribution listed at the beginning of this section. Ordinarily, on the basis of any number of apparently correct[73] perceptual judgements or reports, or even gaze orientation, we might readily attribute this sort of knowledge. However, as the corresponding beliefs will always be false, these knowledge attributions can only be accommodated under a factivoid account of knowledge. We can further emphasize the need for such an account by making the same two observations as we did for visuomotor noise. First, I think that even when we note that these beliefs are noisy, we still judge that these beliefs are known. Even if S's representation of an object O's location is slightly inaccurate, S still knows where O is. Second, even if this judgement proves contentious, we need to grant knowledge in these cases to avoid the threat of skepticism. This skeptical problem becomes clearer over the course of this chapter, as we begin to unravel the full extent of noisy belief content.

In short, we might see that from the mental content that guides perceptual judgement, we might construct a very parallel argument against factivity to that presented in chapter three. The key difference here is that we cannot appeal to the totally inaccessibility of this content to bolster the argument's plausibility. Nevertheless, it seems we might still say that the error is itself inaccessible even if the content generally is. This means that we never have true beliefs about the spatial properties of the objects in our environment, whether we are acting on them or making perceptual judgements about them. Therefore, all the knowledge attributions associated with this perceptual judgement and motor action can only be accommodated by a factivoid theory of knowledge.

3. Proprioception

Even in the absence of visual feedback, we ordinarily think that we have knowledge of the location and orientation of our limbs. We readily attribute knowledge to S of where here limbs are, that they are oriented thus-and-so, etc., even when they are outside of S's visual field. While we might do so on the basis of S's action or verbal reports, upon closer inspection it seems we are willing to attribute such knowledge even in the absence of any explicit behavioural cues. This might be explained by the fact that we always have a sense of limb location and orientation, referred to as "proprioception," which is independent from other sensory input (for a primer see Tuthill & Azim 2018). However, as I will discuss here, proprioception is not perfectly accurate. The representations of limb location and orientation it produces are not perfectly accurate, and in fact systematically biased. This means that if we gain

[73] Of course, "apparently correct" is quite a loaded term here. What I'm trying to convey is that the ordinary language, and even conscious awareness itself, is too coarse-grained to capture the misrepresentation we're interested in here.

knowledge of limb location and orientation on the basis of proprioception, then this knowledge cannot be explained by a factive account of knowledge. Here, however, we need to be careful, as proprioception is famously rudimentary and significantly less accurate than visual perception. This opens the door for a very plausible skepticism of putative proprioceptual knowledge, a possibility I'll leave open here. Thus, in contrast with the other cases I've discussed, for proprioception I'll limit myself to a weaker, conditional claim: If we acquire knowledge of limb location and orientation on the basis of proprioception, then this knowledge violates factivity. In arguing for this claim, I'll first provide a brief background on proprioception, before discussing the associated belief and knowledge attributions.

There are multiple ways in which of our cognitive systems might acquire information about the location and orientation of our limbs. One of these sources is of course vision, and it has also been argued that forward modelling, a predictive capacity for objects in motion, plays a key role in the representation of one's own limbs while they are in motion (Yavari et al. 2015). However, even when our limbs are neither within our visual field nor in motion, we are still able to sense where they are and how they are oriented. This ability is known as "proprioception." While there is some philosophical disagreement whether this constitutes a perceptual capacity (see Gallagher 2003 for an argument that it is not, and Fridland 2011 for a response), on the empirical side it is certainly treated as such. For example, consider the following (emphasis all mine):

> Our ability to properly move and react in different situations is largely dependent on our ***perception*** of our limbs' position. At least three sources — vision, proprioception, and internal forward models (FMs) — seem to contribute to this ***perception***. (Yavari et al. 2015, 403)

> In "Experiment 2", we examined the relationship between actual hand position and the conscious ***perception*** of hand position by subjects during tracking. It was expected that withdrawal of visual inputs would increase the variability of the ***perception*** of dynamic hand positions, and that this may, in turn, account for poorer unseen target position judgements. (Tanaka et al. 2009, 376)

> In the absence of visual feedback, the ***perceived*** position of the hands is systematically biased towards the plausible manual task space. (Fraser and Harris 2016, 3565)

The key here is that the representations produced by proprioception play similar roles to those produced by visual perception. Just as unconscious vision represents the location of objects being acted on, unconscious proprioception represents the position of hands during motor action (Vindras et al. 1998; Jones et al. 2010; Veilleux & Proteau 2011). Similarly, just as conscious vision represents object location and orientation in a way that might guide perceptual judgements, we might also make judgements about limb position on the basis of proprioception alone (Vindras et al. 1998; Plooy et al. 1998). In short, structurally proprioception is very much like visual perception, contributing both to both accessible perceptual content and content for motor action. Accordingly, here I will discuss proprioception simply as a mode of perception.

Like the perceptual representations discussed previously, those produced by proprioception are not perfectly accurate. Behavioural evidence suggests that proprioceptive representations of hand position are systematically biased. As an example, let's look at the widely cited study conducted by Vindras et al., which notes that the bias displayed for reaching in certain conditions suggests that proprioception produces biased representations of hand position (1998). During reaching tasks under suppressed visual feedback, our reaches tend to display a constant error, a "parallel shift" between target endpoint and reaching endpoint (1998, 3290). As discussed in chapter two, specific patterns of error are indicative of specific sources. In brief, here, as the endpoint bias is (i) present in both perceptual judgement and motor action and (ii) eliminated by visual feedback, this suggests that the bias is in the proprioceptive representation itself.[74] They conclude:

> 1) Kinesthetic estimation of hand position may be consistently biased. Some of the mechanisms responsible for these biases are always active, irrespective of whether position is estimated overtly (e.g., with a matching paradigm), or covertly as part of the motor planning for aimed movements. 2) Pointing errors reflect to a significant extent the erroneous estimation of initial hand position. (ibid, 3290)

Similar results have been reported by many other studies (Desmurget et al. 2000; Farrer et al. 2003; Cressman & Henriques 2010; Saidi et al. 2012; Darainy et al. 2013). Yarvari et al. nicely characterize these sorts of findings, observing after one experiment that when only able to access proprioception, "subjects felt their hand's position on the right and bottom of its actual position" (2015, 415). Beyond hand localization, although the subject of less research interest, there is also evidence for bias in the proprioceptive representation of hand orientation (Fraser and Harris 2016/ 2017).

All this means that any beliefs of limb position or orientation whose content is determined by proprioceptive representations will be false. Accordingly, on a factive account of knowledge, we cannot on the basis of proprioception know where our limbs are or how they are orientated. Again, I think we need to leave open the possibility that proprioception is simply not accurate to actually produce knowledge. However, it is clear that if we want to say that we have this knowledge on the basis of orientation, then this can only happen under a factivoid account of knowledge.

[74] As I understand it, the rationale here goes something like this: (i) rules out the possibility it's execution noise, since the bias is in perceptual judgements too. (i) also rules out the possibility that it's something specific to motor control or conscious perception. (ii) rules out coordinate transformations or some higher-level source for the bias, since this would persist for vision. Thus, given all the this, the best explanation is that the bias is in proprioception itself.

4. Analogue Magnitude Representation

Analogue magnitude representation (AMR) is a specific representational kind, which, as I will argue, is especially problematic for factive accounts of knowledge. In this section, I will first introduce the concept of AMR, after which I will discuss two specific instances, analogue *numerical* magnitude representation and analogue *temporal* magnitude representation. Both cases present a clear argument against the factivity of knowledge. Ultimately, if we wish to retain our ordinary knowledge attributions involving numerical and temporal magnitude, we need to adopt a factivoid framework.

4.1 Introduction to AMR

Much of the information our cognitive systems process assumes the form of magnitudes. Ordinary tasks will often involve such information as the quantity of a given entity (numerical magnitude), the size of an object (spatial magnitude[75]), or elapsed time (temporal magnitude). Accordingly, the successful completion of ordinary tasks will often require accurate representations of magnitudes, representations that are often closely associated with judgement and belief. For example, the judgement that one group of objects contains more members than another group requires the accurate representation of the numerical magnitude of membership of each group. Moreover, on the basis of S's judgement that group A contains more members than group B, it is ordinary practice to attribute to S beliefs about the magnitudes in involved in the judgements. Finally, under ordinary conditions, we attribute knowledge to S if indeed group A contains more members than group B, such as knowledge of the magnitude of membership of each group. This of course is not unique to numerical magnitude, and we might say something similar for judgements and beliefs involving temporal magnitude.

The importance of magnitude representation in the guidance of behaviour raises the question of how magnitude is represented internally. Put simply, there is widespread agreement that magnitude representation is **analogue**. Not only is this the case in the vast empirical literature (for overviews see Ansari Vogel 2013; Piazza & Eger 2016) but also in what minor attention it has received on the philosophical side (Beck 2015; Gemel & Quinon 2015). Before surveying the evidence for general analogue magnitude representation, however, I should first say a bit on what it means for a representation to be analogue.

At the most basic level, an analogue representation is exactly what its name might suggest: A representation is a direct analogue for the information being represented. This idea is best conveyed via example, likely the most familiar of which is not neural, but rather that of analogue audio signals. In the case of an analogue audio signal, voltage as a function of time directly represents sound pressure as a function of time. Higher sound pressure is represented by higher signal voltage, and lower pressure by lower voltage, via a continuous function. In this way, we might understand voltage to represent sound pressure as its direct

[75] Of course, I have opted to discuss size in conjunction with other spatial properties that guide motor action (i.e. location and orientation) and thus won't retread it here.

representational analogue. We might contrast this with digital audio signals, on which voltage is not a direct analogue for sound pressure, but rather corresponds with a binary code, which in turn corresponds with a certain level of sound pressure. The key here is that for digital representation, there is no direct analogue for sound pressure, nothing that represents increasing sound pressure by simply increasing, and so forth.

Just as voltage is a direct analogue for sound pressure in analogue audio representation, magnitude is represented by an analogue in the brain, "some neural entity that is a direct analogue of the magnitude it represents" (Beck 2015, 835). In order to illustrate what this might look like, let's consider the example of the "mental number line." While I'll discuss the details more in the following subsection (see Ansari Vogel 2013 for an overview), the basic idea is that numerical magnitude is represented by a neural structure that functions like a sort of "number line," with any given numerical magnitude represented as a specific location along this number line.

The general evidence that neural magnitude representation is analogue is both behavioural and neurological. While in the following subsections I'll highlight specific evidence pertaining to numerical and temporal magnitude, here I want to briefly highlight the thrust of the behavioural evidence for AMR, which involves two key observations: (1) Magnitude representation occurs independently from semantic representation, and (2) magnitude representation corresponds with Weber's law. I'll discuss each in turn.

There is a wealth of behavioural evidence that highlights that magnitude representation is not contingent upon the capacity to process written language. First, we might observe that pre-linguistic infants demonstrate the capacity to accurately represent magnitudes. For example, a study by vanMarle and Wynn reported that "by the age of 6 months, infants have a means of representing duration, and they can discriminate durations of audiovisual events" (2006, F44). This conclusion is based on the observation of infants fixating much longer on visual targets (in this case a puppet) when it emits a tone of a duration discriminable from the previous tone emitted; in this case the discrimination threshold was 2:1 (ibid). Similar magnitude representation in infants has been reported in a multitude of other studies, including for numerical magnitude (Xu & Spelke 2000; Lipton & Spelke 2003; Xu et al. 2005; vanMarle & Wynn 2009), spatial magnitude (Brannon et al. 2006), and velocity (Daum et al. 2016; Möhring et al. 2017) in addition to temporal magnitude (Brannon et al. 2007; Hevia et al. 2017).[76] Beyond this, there is also evidence of magnitude representation in non-human primates (Beran et al. 1998; Cantlon et al. 2007; Tomonaga 2008; Jordan et al. 2008; Gazes et al. 2017) and even birds (Pepperberg & Carey 2012; Rugani et al. 2015a/b).[77] All this is important because, although it doesn't definitively answer *how* magnitude is represented in adult humans, it suggests that magnitude need not be represented symbolically (see Dehaene et al. 1999). That is, it suggests, for example, we might represent the magnitude of 12 independently from the

[76] It is so thoroughly established that infants represent magnitude that current debate features such questions as whether they *exclusively* use AMR (Mou & Vannarle 2014) and whether magnitude is represented in a single system (Lourenco & Longo 2010).
[77] For a review see Davis & Pérusse (1988).

"semantic representation of numerical symbols" (Piazza & Eger 2016, 264) like "12." Crucially, this suggests that our behaviour in judgements of comparison might be explained not by any semantic representations, but by non-semantic representations of magnitude, which we have both general and specific reason to assume are analogue.

The most compelling behavioural evidence that magnitude representation is analogue emerges from the fact that magnitude representations follow a pattern referred to as "Weber's law," on which "the ability to discriminate two magnitudes is determined by their ratio" (Beck 2015, 833). For example, in a number discrimination task, we might note that we are faster to judge that 2 is larger than 1 than we are that 9 is larger than 8 (Ansari and Vogel 2013). Despite the fact that there is a difference of one in both cases, 2:1 is larger than 9:8. Weber's law is simply this correlation of discrimination of magnitude with the ratio of the magnitudes being discriminated. This basic pattern holds for the above cases of infants and non-human animals, as well as adult humans, as I'll discuss in detail in the following subsections. For example, in the vanMarle and Wynn study discussed above, infants were able to discriminate between temporal magnitudes for a 2:1 ratio, but not a 3:2 ratio (2006).

The reason this constitutes evidence that magnitude representation is analogue reduces to one key observation: On an analogue representational format, as magnitudes get larger their representational forms become more similar. This means that the representations of two large magnitudes of difference X are more similar than the representations of smaller magnitudes of difference X, and thus require additional cognitive resources to discriminate. An excellent way to understand this is with Carey's illustration (2009, 118), which I'll recapitulate as the following: Say instead of representing numerical magnitude with numerals, we represent magnitude in an analogue format, for which one "-" corresponds with a single numerical unit. That is "-" represents the magnitude of 1, "--" the magnitude of 2, and so on. Here we note that as magnitude increases these representations become more similar, in that they have more in common. 9 and 8 (--------- and --------) are far more similar in representational form than 2 and 1 (-- and -), despite sharing the same difference. Thus, we might note it is much more difficult to visually discriminate 9 and 8. Conversely, 8 and 4 (-------- and ----) are just as discriminable and 2 and 1 (-- and -), as they have the same ratio. As I'll discuss with numerical magnitude specifically, there is additional neurological evidence for analogue magnitude representation. However, here we might conclude that the fact that magnitude representation follows Weber's law is easily explained if magnitude representation is analogue, a fact which has led to the widespread acceptance of the claim that magnitude representation is analogue (Brannon 2006; Spelke & Kinzler 2007; Carey 2009; Beran et al. 2011; Beck 2015).

While I'll discuss specifics in the following subsections, as they pertain to numerical and temporal magnitude respectively, here we might note the general concern that AMR raises for the factivity of knowledge. Put simply, analogue representations tend to be noisy. Even before focusing attention on neural analogue representation, analogue audio signals foreshadow this problem. The familiar ubiquity of noise in analogue audio signals is not a quirk of audio transmission, but rather a basic feature of analogue representation. Although perhaps obscured by the above toy example, on which the analogue of numerical magnitude was a set of discrete dashes, actual analogues tend to be continuous, such as voltage distributions for analogue audio signals. Thus, slight different voltages, for example, result in the representation of slightly different audio pressures. This means that any fluctuation in the signal (i.e. noise) changes what is represented. In this way, input noise is transmitted through to the output. Since input noise is unavoidable, so is output noise. As we will see in §4.2 and §4.3, the continuous neural distributions used as analogues for the representation of numerical and temporal magnitude both succumb to this problem.

This poses the same challenge for factivity as the other cases of noise discussed previously. If our representations of magnitude are noisy, then this means that they don't represent the world as it is and, more crucially, that they slightly misrepresent the world. We might make a similar argument that these representations determine belief content, on the basis that they explain the behaviour on which we form certain belief attributions, and note that, given the belief still meets its output requirements, we still judge this belief to be known. However, as the content of these beliefs misrepresents the world, they are false. Therefore, we need a factivoid account to handle them. Again, note that this is only a sketch of the general argument, and I'll fill in specifics for each type of magnitude representation discussed below.

The final thing I want to note before moving on to the specifics of numerical and temporal magnitude is that while it is widely agreed adult humans have a cognitive system that represents magnitude via a neural analogue, it is not widely agreed that this is the *exclusive* means by which we might represent magnitude. For example, it has been proposed that numerical magnitudes between one and four might be represented by a separate "object tracking system" (Mou & VanMarle 2014; see also Feigenson et al. 2004; Revkin et al. 2008; Trick 2008).[78] This underscores the importance of establishing that the belief content I'll discuss in the following subsections is in fact determined by AMR, which is a matter of demonstrating that the behaviour on the basis of which we attribute relevant beliefs is clearly guided by the analogue representation of numerical (or temporal) magnitude.

4.2 Numerical Magnitude

In this subsection, I'll argue along the following lines that much of our knowledge involving numerical magnitude will violate factivity: Due to the analogue way in which the brain represents numerical magnitude, it appears that many of our beliefs involving numerical magnitude will in fact be false. In certain cases we might judge that these beliefs are simply unknown, especially when involving large numbers. However, in more everyday cases, as with the false beliefs that guide motor action, we still judge that these beliefs are known. Beyond this, in order to avoid skepticism for some knowledge involving numerical magnitude we must grant that these false beliefs are often still known. As such beliefs are still highly truthlike, meet the output requirements of the cognitive processes according to which they

[78] Carey argues for something similar (2009, ch. 4).

form, still play the same functional role as true belief, etc., they fit perfectly into the factivoid framework outlined in the previous chapter. In order to present this argument, this subsection will be structured in the following way. First, I'll provide an overview of the empirical evidence that numerical magnitude specifically is represented via a neural analogue, a sort of mental number line. Here I'll focus on evidence form behaviour and neuroimaging.[79] From this, I will argue that much of our knowledge involving numerical magnitude is factivity-violating. We see this especially with attributions like "S knows what the magnitude of n is."

As we've seen above, numerical magnitude tends to be the paradigm example of AMR, and it is widely accepted that (at least) the vast majority of numerical magnitude representation is accomplished internally via a neural analogue. This neural analogue is often referred to as a "mental number line," as numerical magnitude is represented topographically, left-to-right, much as it is on a physical number line (Göbel 2006; Harvey et al. 2013). The specifics of this have been the subject of debate[80]. However, the standard view is that numerical magnitude representation is not represented by discrete unit (neuron) activation, but by an activation distribution.[81] Consider the following account of how, after processing, numerical magnitude is finally represented:

> In the final stage, input is received from the summation code module and the summed activity[82] is projected onto a topographically organized "number line." In other words, the summed activity is mapped onto a selective numerical magnitude unit that only responds to a selective range of values, but not if the value is larger or smaller. For example, the summed activity of seven objects is projected onto a unit, which prefers the numerical magnitude 7. However, this unit will also respond to the numerical magnitude 6 and 8, although to a lesser degree. Therefore, each projected numerical magnitude creates a bell-curved Gaussian activity distribution where the peak activation is centered on the preferred numerical magnitude. (Ansari and Vogel 2013, 394)

In short, numerical magnitude representation is analogue. The magnitude of any given number is represented by a Gaussian distribution. Before discussing the implications of this type of representation, I first want to highlight the empirical evidence for the claim that numerical magnitude is represented via a continuous neural analogue. This comes both from the behavioural evidence that numerical magnitude representation follows Weber's law, which suggests the representation is analogue, and neuroimaging evidence that also suggests that it is analogue.[83]

---

[79] I must note again that much of the evidence cited in the literature on numerical magnitude comes from neurophysiological studies on monkeys. As with the "two visual systems" hypothesis in chapter two (§3), due to the horrifying ethical implications I want to minimize the role this research plays in my account.

[80] The key disagreement here is between whether representation is logarithmic (Dehaene & Changeux 1993) or linear (Gallistel & Gelman 2000). For a general overview of this distinction see Ansari and Vogel (2013).

[81] Note that a rival account has been advanced by Zorzi et al. (2011).

[82] "Summation code module" and "summed activity" refer to the processing stage prior to final magnitude representation, which I discuss more shortly.

[83] For a more sophisticated review of this literature, see Piazza and Eger (2016).

Beginning with the findings of the seminal 1967 study by Moyer and Landauer, it has consistently been observed that the time required to judge which of two numbers is larger in magnitude is (roughly) inversely proportional to their ratio. That is, numerical magnitude representation follows Weber's law. This result has been confirmed in the results of a multitude of subsequent studies (Parkman 1971; Buckley & Gillman 1974; Banks et al. 1976; Gibbon 1977; Dehaene & Akhavein 1995). Put one way, "the most consistently observed effect in numerical cognition research is that the difficulty of numerical discrimination depends on the ratio difference between two values" (Prather 2014, 601). As discussed above, this strongly suggests that numerical magnitude representation is analogue.

Beyond this, evidence from neuroimaging studies also supports the claim that the representation of numerical magnitude is analogue. Before discussing this, however, I first want to note that neurological study of numerical representation is complicated both by the fact that (i) numerical magnitude might be represented via multiple external formats (ex. "2" vs. "two" vs. *two* vs. * *) (Ansari 2007; for a meta-analysis see Sokolowski 2017a) and (ii) multiple regions of the brain contribute to the processing of numerical magnitude (for a meta-analysis see Sokolowski et al. 2017b). This results in a complex picture of numerical magnitude representation at the neural level, the details of which are still the subject of present research and debate (Sokolowski et al. 2017a/b). Rather than delving into the details of these complications, I'll instead opt to simply highlight the fact that when a neuroimaging study observes an apparent analogue representation for a numerical magnitude, this does not necessarily mean that all numerical magnitude is represented in this way, or even that that specific magnitude is not also represented another way in the brain concurrently. Again, as first discussed in chapter two, it is important to keep in mind that the same information is often represented internally multiple times, in multiple ways.[84] With this in mind, we might approach the evidence from neuroimaging studies that numerical magnitude representation often is analogue.

Perhaps the most obvious neuroimaging evidence for analogue numerical magnitude representation is the observation that activation increases monotonically as the magnitude of a set of objects in a visual field increases. That is, "the more items present of a given non-symbolic stimuli, the higher the neurofunctional activation within the posterior parietal lobe" (Ansari & Vogel 2013, 395). For example, a study by Satens et al. (2010) presented participants with visual arrays that contained between one and five dots and used fMRI to measure corresponding activation. They found "a positive correlation between the strength of the BOLD[85] response and number of target locations" (ibid, 82). As activation increases with magnitude, this of course is indicative of an analogue representation, with activation levels in this level of the brain serving as a neural analogue for magnitude. Similar results have been observed in other neuroimaging studies (Roggeman et al. 2010; Roggeman et al. 2011; Park et al. 2016). This particular activation pattern is often cited as evidence for "summation coding," an intermediary stage of magnitude processing on which the magnitude of a

[84] This underscores the importance of behavioral evidence in establishing which representations actually determine belief content.

[85] ***B***lood ***O***xygen ***L***evel ***D***ependent, a neuroimaging technique for observing brain activity

collection of items is simply summed before being mapped onto the mental number line (Ansari & Vogel 2013, 395; Sokolowski 2017b, 390).

There is also neuroimaging evidence for final analogue representation of magnitude on the mental number line, which at the neural level is often referred to as "place coding" (Ansari & Vogel 2013; Sokolowski 2017b, 390). An early neuroimaging study to confirm this was conducted by Piazza et al. (2004). Researchers used fMRI to observe neural activation when participants viewed visual arrays containing either 16 or 32 dots, finding that:

> Although noisier than behavioral curves, the brain activation responses to habituation numerosities 16 and 32 showed the characteristic features of Weber's law. The curves had an increasing width on a linear number scale, became identical in width when plotted on a log scale and could be expressed as a simple Gaussian function. (ibid, 550-1)

In short, these representations displayed all the hallmarks of being analogue, even at the neural level. Similar conclusions have been reported by a number of subsequent studies (Piazza et al. 2007; Jacob and Nieder 2009; Eger et al. 2009; Harvey et al. 2013).

As discussed in §4.1 with respect to AMR generally, if numerical magnitude representation is analogue, then it will be noisy. Thus, because of the behavioural and neuroimaging evidence I've just overviewed, it is widely accepted that "representation of numerical magnitude is analogue and, therefore, inherently noisy and approximate" (Ansari & Vogel 2013, 397). This is the case both for the significant empirical (Gallistel & Gelman 2000; Piazza 2010; Prather 2014; Nieder 2016) and the much smaller philosophical (Beck 2015; Égré 2017) discussion[86] of numerical magnitude representation. As with visuomotor noise, it is worth noting that noise in the representation of numerical magnitude does not necessarily undermine the success of the behavior it guides. For example, useful especially for magnitude discrimination tasks, "this system yields a noisy representation of approximate number that captures the inter-relations between different numerosities" (Feigenson et al. 2004, 309). Just as the slight misrepresentation of an object's location does not hinder our ability to successfully grasp that object, the slight misrepresentation of two magnitudes does not hinder our ability to judge which magnitude is larger.

Now I would like to begin the turn from empirical findings to the epistemological problem posed by these findings, the core of which should now sound familiar: Analogue representations of numerical magnitude are inaccurate in the sense that numerical magnitude is systematically misrepresented. For AMR of numerical magnitude, there are three distinct sources of systemic misrepresentation, which I discuss in turn: (i) bias, (ii) noise, (iii) and approximation. While the first two of these will of course be familiar from previous discussions, the third is unique to AMR. Accordingly, I will afford it the most attention.

As I'll use the term here, analogue representations of numerical magnitude are "biased" in the sense that magnitude of certain numbers are be consistently

[86] Much of this discussion seems to focus around vagueness in the philosophy of language (Fults 2011; Gemel & Quinon 2015; Égré 2017).

misrepresented relative to other magnitudes. For example, if one's internal representation of the magnitude of 5 happens to be consistently closer to the magnitude of 9 than it is to 1, this sort of systematic overrepresentation of the relative numerical magnitude of 5 is a bias in the representation of the magnitude of 5. While the observed patterns of relative representation of numerical magnitudes tend to be complex[87], and it is not universally agreed upon how these representations are biased[88], the prevailing logarithmic account maintains "a relative overrepresentation of small numbers compared to larger numbers" (Ansari & Vogel 2013, 394). Additionally, an especially robust effect is the systematic misrepresentation of the midpoint of two numbers (Göbel et al. 2006; Longo & Lourenco 2007; Loftus et al. 2008; Mohr 2010). Although not the most important contribution to inaccuracy I'll discuss, it is at least worth noting apparent bias in the analogue representation of numerical magnitude.

Next, there's noise. To the extent that numerical magnitude is represented via neural analogue, those representations will be noisy. The argument from this point unfolds similarly to that from motor action, the first step of which is to argue that this noise can be understood not only as indeterminacy but also inaccuracy. That is, the noise in our representation of numerical magnitude means that we systematically misrepresent numerical magnitude. Or, as put by Égré in one of the few philosophical papers to directly address the matter (emphasis mine):

> An example of analog measurement is given by the approximate number system, by which we estimate numerosities. A key feature of this system is that cardinalities are represented with some probabilistic error, ***meaning that the same number of objects will be represented as more or less than it is***. (2017, 3838)

What Égré is referring to here is trial-to-trial variability in the represented magnitude of a set number of objects. That is, for certain trials magnitude will be overrepresented, and for other trials it will be underrepresented. We might compare this trail-to-trial variability to that we might observe for motor action. Just as our reaching and grasping movements are variable as the result of noise in our representations of relevant spatial properties, our judgements about the magnitudes of numbers vary as the result of noise in our representation of numerical magnitude.[89] This is best illustrated in the trial-to-trail variability of responses in number line tasks. Pioneered by Siegler & Opfer (2003), it is now widely accepted practice that we can observe subjects' representations of numerical magnitude by having them indicate those magnitudes on number lines (e.g. Siegler & Booth 2004; Opfer & DeVries 2007; Geary et al. 2008; Berteletti et al. 2010).[90] It is well documented that responses on these tasks are variable (for

---

[87] For example, a series of studies by Landy et al. (2013; 2017) documented a complex pattern of relative misrepresentation involving large numbers, including discontinuities and, most notably, two distinct populations of participants.

[88] See Gallistel and Gelman (2000) for the alternate, "linear" account, and Karolis et al. (2011) for evidence for bias corresponding with that account.

[89] We might even note that variability in motor action also displays a Gaussian distribution (Bock & Daunicht 1987).

[90] For an alternative account, see Barth and Paladino (2011).

studies that highlight this, see Cohen & Blanc-Goldhammer 2011; Kim and Opfer 2017; Peeters et al. 2017), which suggests trial-to-trail variability in the representation of magnitude. In short, as Égré put it, numerical magnitude "will be represented as more or less than it is" (2017, 3838). This variability derives from the noisy nature of individual neural activation. Per Nieder:

> The read-outs for single trials of neuronal responses are noisy, and smooth and unequivocal tuning functions only emerge after the discharges of a cell to many identical trial repetitions have been averaged, or after the tuning functions of many equally tuned single cells have been pooled. (2016, 371)

It is uncontroversial that the activity of individual neurons is noisy (see Faisal et al. 2008).

Finally, there is approximation. While it is widely accepted that analogue representations of numerical magnitude are approximate, it is a matter of debate what exactly this means (see Beck 2015, §3.2). Here I will not try to address this matter generally, but only highlight one specific point of the approximate representation of numerical magnitude: The neural analogue for numerical magnitude represents multiple magnitudes simultaneously. That is, we might say that the same number of objects will be represented as more *and* less than it is. As Nieder continues from above (emphasis mine):

> The brain, however, needs to decipher a presented number within a single trial. To cope with noisy discharge patterns, it must take the activity of a population of other tuned neurons into account. The important information in a population code is thus the relative, rather than the absolute, amounts of activity in different neurons. Just like a democratic election involves many voters, the coding of a particular number may involve many neurons, and ***the signal of each number neuron represents an independent vote for a specific number***. (2016, 371)

What Nieder is illustrating here is that since the representational content of individual neurons is simply whatever fact about the world they're tuned to (see deCharmes and Zador 2000), whenever a neuron tuned to number *n* discharges, that is taken to represent the magnitude of *n*. However, crucially, the AMR content that guides behaviour involving *n* will not only be comprised of neurons whose activation represents *n*, but also neurons whose activation represents magnitudes greater than *n* and less than *n*. Thus, while this content represents the magnitude of "*n*" as *n*, it also represents it as more than and less than *n*. In this manner we can understand the fundamental misrepresentation inherent in the analogue representations of numerical magnitude. This content "approximates" the magnitude of *n* by representing it as *n*, but also to a lesser extent *n+1, n-1,* and so forth. Therefore, as this content always misrepresents the magnitude of *n* as being both greater than and less than *n,* beliefs that have content determined by these analogue representations will be false.

Before continuing, I want to briefly address two potential worries: (1) Why don't we just say that the behaviour-guiding content is determined only by neurons displaying the greatest activity? (2) Why don't we just say that individual neurons don't represent specific magnitudes, only activation distributions do? The answer to (1) is simply that, as discussed at length above, doing so leaves us unable to explain behavioural observations like Weber's law. Based on our behaviour, it is clear that the AMR content that guides magnitude comparison tasks, etc., includes the entire set of active neurons on the neural number line, not just the one that is most active. The answer to (2) is that taking this stand would be to reject the notion of neural representation entirely. While the exact nature of neural representation is not widely agreed upon, it widely accepted that individual neurons do have their own representational content (for a review see Vilarroya 2017).

In the end, due to the biased, noisy, and approximate nature of neural analogue for numerical magnitude, we can understand that these representations are inherently inaccurate. Of course, the fact that analogue representations of numerical magnitude are inaccurate provides a case for factivoid knowledge only if our beliefs contain AMR content. At first glance, beliefs involving vague predicates (e.g. "tall" or "approximately 100") seem like a promising avenue, as this is within philosophical precedent. For example, Gemel & Quinon maintain that analogue representations of numerical magnitude provide "the correct semantics for exact natural numbers, especially when used as vague quantifiers" (2015, 91). However, vague predicates turn out to be a bit of a dead end, as it's not obvious that AMR in these cases actually misrepresents vague predicates. It's not clear, for example, that when we misrepresent S's height this means that we also misrepresent that S is tall. This I think is a complex and interesting question, as I indicate in §5. However, for the purposes of this specific argument, the better option is to look for beliefs with more obviously misrepresentative content. In order to do this, I want to start from the behavioural level, focusing on the behaviour that is guided by analogue representations of numerical magnitude. After identifying this behaviour, we might ask what beliefs we attribute on the basis of this behaviour, whether these beliefs are false, and whether these beliefs are known.

As already discussed above, because they follow Weber's law, it is widely accepted that numerical magnitude discrimination tasks utilize analogue representations. Moreover, there is compelling evidence suggesting that both symbolic and non-symbolic arithmetic tasks are also supported by analogue representations. Neuroimaging suggests that the neural substrates of AMR are shared by those of arithmetic (see Knops & Willmes 2014) and, behavioural observation from developmental psychology indicates that development of numerical AMR correlates with acuity in arithmetical tasks (Halberda et al. 2008/2012; Piazza et al. 2010; Lyons & Beilock 2011). In short, we can safely assume that our behaviour in both numerical discrimination and arithmetical tasks are guided by analogue representations of numerical magnitude.

Rather than focusing on the beliefs we might attribute on the basis of these behaviours, I want to skip directly to associated knowledge attributions. The reason for this is that language is limited in its ability to describe these beliefs, and it's easier to get to them if we start with knowledge attributions and work backwards. Under standard, everyday conditions in which S displays a reliable ability to make numerical judgements guided by AMRs, it is ordinary practice to attribute to S knowledge of the magnitudes involved. For example, if S judges that n is less than n+1, we might say that (i) S knows the magnitude of n, (ii) S knows what the magnitude of n is, or (iii) S knows that the magnitude of n is "n." However, a complication arises for these knowledge attributions when we consider that our representations of numerical magnitudes are inaccurate. This means that "n," S's analogue representation of the magnitude of n, slightly misrepresents the magnitude of n. Taking the definite description "the magnitude of n" to refer to the actual magnitude of n, and ""n"" to refer to S's internal representation of the magnitude of n, this means that S's belief that the magnitude of n is "n" is false. Again, this is because S does not possess a perfectly accurate representation of the magnitude of n. This means that on a factive account of knowledge, S cannot know what the magnitudes of numbers are.[91]

It is my contention that, even when we note that our representations of numerical magnitude are biased, noisy, and approximate, we still judge that we can have knowledge of numerical magnitudes. After all, these representations still meet the output requirements for the cognitive process that produces them. As mentioned above, these representations still retain the interrelations of the magnitudes they represent, and are sufficiently accurate to guide reliable judgements (at least for relatively small numbers[92]). Here we have another case in which it is only through empirical research that we might even realize that these representations are systematically misrepresentative. Taking this all into account, I don't think the experimental evidence presented here changes our intuitive judgements that we have plenty of knowledge about the magnitudes of numbers. However, again I recognize that this judgement might not be entirely universal, in which case we might again appeal to the ever-growing skeptical problem associated with the factivity of knowledge. From previous chapters and §2, we might note that retaining factivity means conceding that we don't ordinarily know where objects are in our visual field, how large those objects are, how they are orientated. Now, we can see this also means conceding that we don't know what the magnitudes of numbers are. Accordingly, assuming we wish to avoid the skeptical conclusion, we must adopt a factivoid theory of knowledge, as discussed in the previous chapter.

In this subsection, I've detailed evidence for the analogue representation of numerical magnitude, the properties of these representations, and the epistemological implications of these properties. Crucially, we observed that these representations are biased, noisy, and approximate, all features that contribute to systematic misrepresentation of numerical magnitude. However, we might also observe there are knowledge attributions associated with the behaviour guided by this content. Ordinarily, we want to say that S knows what the magnitudes of numbers are. This means that in order to preserve this attributive practice and avoid a skeptical result for knowledge of numerical magnitude, we need to adopt a factivoid account of knowledge

[91] I should note here that, as discussed above, there is also evidence for a second cognitive system used to represent magnitudes 1-4. Thus, we must restrict the domain of this conclusion to magnitudes greater than 4.

[92] Our representation of the magnitudes of large numbers (>100,000) might be inaccurate enough that we don't want to consider them to be knowledge (see Landy et al. 2013; 2017).

4.3 Temporal Magnitude

I want to finish my discussion of AMR by saying a bit on the representation of temporal magnitude. Rather that discussing it with same depth I did with numerical magnitude, here I only want to highlight two empirical points: (1) Representations of temporal magnitude appear to largely share the same mechanisms as representations of numerical magnitude, and (2) the representations are not perfectly accurate. On this basis, we might argue along the familiar lines that much of our knowledge involving temporal magnitudes can only be accommodated by a factivoid account of knowledge.

In a widely cited 2003 paper, Walsh proposes that "time, space and quantity are part of a generalized magnitude system" (483). He writes:

> It is the purpose of this paper to bring together, as A Theory Of Magnitude (ATOM), disparate literatures on time, space and number, and to show similarities between these three domains that are indicative of common processing mechanisms, rooted in our need for information about the spatial and temporal structure of the external world. (ibid, 483)

Here I will specifically focus on the unifying framework for numerical and temporal magnitude representation. While Walsh also based his proposal on both neuropsychological and neurophysiological evidence, neuroimaging is most relevant for linking time and number. Walsh cites multiple studies that indicate common neural substrates for the processing of numbers and time (Chochon et al. 1999; Dehaene et al. 1999; Rao et al. 2001; Piazza et al. 2002). While the extent to which these neural correlates overlap is still not agreed upon, it is recognized that there is at least some shared cognitive resources between numerical and temporal magnitude processing (Bueti & Walsh 2009; Skagerlund et al. 2016).

Beyond this, there is also a wide body of evidence that suggest that temporal magnitude is represented spatially in general (see Leone et al. 2018, §1), as well as specifically in a "mental time line" (Bonato et al. 2012; Di Bono et al. 2012; Marin et al. 2016). This of course we might compare with "mental number line" discussed in §4.2. Finally, there is a good deal of behavioural evidence demonstrating "interference" between representations of temporal and numerical magnitude, on which the processing of one impacts or "interferes" with the processing of the other. One representative study by Xuan et al. found that "stimuli with larger magnitudes in these nontemporal dimensions were judged to be temporally longer" (2007, 1). This means, for example, one might judge the duration of a presentation of a stimulus of 10 dots to be longer than that of a stimulus of 5 dots, presented for an equal duration. Similar results have been reported by a number of subsequent studies (Oliveri et al. 2008; Vicario et al. 2008; Chang et al. 2011; Hayashi et al. 2013). All this suggests that temporal magnitude is represented very similarly to numerical magnitude.

In addition to the general suspicion we might have regarding the accuracy of our representations of temporal magnitude—derived from their shared neurocognitive mechanisms with the biased, noisy, and approximate, representation of numerical magnitude—we might also note there is also more direct evidence that these representations are inaccurate. Most significantly, these representations are biased in a very specific way:

> The relative inaccuracy of temporal judgements in the visual modality appears not simply to comprise greater variance about an accurate mean; rather, a consistent bias has been observed, often expressed as "sounds seem longer than lights". (Resta et al. 2016, 101)

Here Resta et al. are referring specifically to findings of Wearden et al., who observed not only that, as we would expect from noisy representations, judgements about temporal magnitude were randomly variable, but also that "representations ran faster for auditory than for visual stimuli" (1998, 97). That is, participants consistently judged the same temporal magnitudes to be greater if it was indicated by an auditory stimulus than if it was indicated by a visual stimulus. This sort of bias in the representation of temporal magnitude was well established by the time of the Wearden et al. paper (Goldstone et al. 1959a/b; Behar & Bevan 1961; for a review see Goldstone & Lhamon 1972). In short, we might safely conclude that our representations of temporal magnitude are not perfectly accurate.

As evidence for the inaccuracy of these representations come from behavioural studies involving direct judgements of the duration of time, we might also conclude that these inaccurate representations are what guide our judgements about temporal magnitude. From here, we might make the same sort of argument against factivity I made in more detail with respect to the inaccurate representations of numerical magnitude. This time, we might note that on the basis of judgements about durations of time, we might attribution beliefs to S about how long that duration was. If S's beliefs are sufficiently accurate, successfully guiding behaviour and so forth, we might also judge that these beliefs are known: S knows how long that event lasted, etc. However, we can only preserve these knowledge-attributing practices under a factivoid account of knowledge. Again we find another set of knowledge attributions that cannot be accommodated if we wish to retain factivity.

5. Future Generalization

The instances of knowledge this project has focused on have all been what we might call "low-level," in that they involve belief content that is to a significant extent inaccessible. This inaccessibility is reflected in the natural language sentences we use in the attempt to describe this knowledge. S knows that the object is there; S knows the size of object O; S knows what the magnitude of number n is, etc. S cannot access much of the mental content that constitutes this knowledge, and therefore we must attribute that knowledge to her in this vague sort of way. Accordingly, we can imagine that the philosopher with a bit of tolerance for skepticism might be tempted to, in lieu of adopting factivoid knowledge, write off the category of "low-level" knowledge entirely. They might argue that the skeptical conclusion in this case is quite acceptable, as "low-level" knowledge is largely incommunicable from person to person, and therefore not what we normally think of when we first take inventory of our ordinary knowledge.

For the purposes of this project, I have simply assumed that any even moderately skeptical result should be avoided. Accordingly, if one is willing to accept that we ordinarily don't have knowledge of the spatial properties of objects in our visual field, numerical magnitude, etc., this is sufficient to block the argument I've presented in favour of factivoid knowledge. However, here I want to at least briefly consider how we might go about generalizing my argument further, so that it also includes "high-level" knowledge, which involves more accessible content. This means addressing questions like:

> *(1) If S doesn't know where the cup is, does S know that the cup is on the table?*
>
> *(2) If S doesn't know the magnitude of 2, does S know that 2+2=4?*

To be clear, I don't have a good answer for these types of question, and here I won't attempt to provide one. However, here I at least want to point out that questions like (1) are a bit misleading, as our representations of the qualitative spatial relations of objects are distinct from the quantitative representation of their location (see Laeng et al. 2003). That is, the content of our belief that object X is in qualitative spatial relation R with object Y is not computed by "taking the quantitative representation and making it coarser" (Laeng 2014, 41). Thus, to determine whether there is any misrepresentation in the content of these beliefs, we more or less need to start from scratch. We might contrast this with (2). As discussed in the previous section, arithmetic is guided by analogue representations of magnitude. Ultimately, in either case, this is a matter for future research, and in this way we might understand that the full extent to which this project generalized is currently an open question.

Conclusion

In this project, I have argued against the factivity of knowledge and developed the non-factive analysis of knowledge, largely on the basis of empirical findings. Taking a step back from this endeavour, I think the most important lesson to take away here is not anything specifically to do with the factivity of knowledge. Rather, I think the most important contribution of this project is its demonstration of the profound role empirical findings can play in the analysis of knowledge. Without research into planning noise (and "Two Visual Systems," and foveal bias, and AMR, etc.) there simply is no argument against factivity, and therefore no reason for the factivoid analysis of knowledge. Moreover, if we hold a few key epistemological assumptions fixed, we might note that whether knowledge is factive reduces to an empirical question. Of course, the fact that it no longer seems like knowledge is factive is itself a very significant finding. But as I have argued here, this might be easily accommodated by standard epistemology. However, the fact we reach this conclusion on largely empirical grounds highlights just how valuable this research can be, even for the most fundamental, seemingly *a priori* programmes of epistemology. What also strikes me is that all this research was conducted entirely independently of epistemological concerns, and I hesitate to imagine the impact cognitive-neuroscientific methods deployed in the service of epistemology. In the end, the future looks especially bright for the continued and increasing use of empirical findings in epistemology.

References

Abath, A. (2012). Epistemic Contextualism, Semantic Blindness and Content Unawareness. *Australasian Journal of Philosophy,90*(3), 593-597.

Achenbach, J. (2005). Cat Carrier: Your cat could make you crazy. National Geographic 208.

Ackermann, R. (1973). *Belief and knowledge* (Problems in philosophy). London: Macmillan.

Adam, J., Davelaar, J., Gouw, E., & Willems, A. (2008). Evidence for attentional processing in spatial localization. Psychological Research, 72(4), 433-442.

Adams, M. (2009). Empirical evidence and the knowledge-that/knowledge-how distinction. *Synthese, 170*(1), 97-114.

Adams, W., Gray, K., Garner, M., & Graf, E. (2010). High-level face adaptation without awareness. *Psychological Science, 21*(2), 205-10.

Aglioti, S., Desouza, J., & Goodale, M. (1995). Size-contrast illusions deceive the eye but not the hand. *Current Biology, 5*(6), 679-685.

Alameda, J., Cuetos, F., & Brysbaert, M. (2003). The number 747 is named faster after seeing Boeing than after seeing Levi's: Associative priming in the processing of multidigit Arabic numerals. *QUARTERLY JOURNAL OF EXPERIMENTAL PSYCHOLOGY SECTION A-HUMAN EXPERIMENTAL PSYCHOLOGY.*

Amitay, S., Guiraud, J., Sohoglu, E., Zobay, O., Edmonds, B., Zhang, Y., . . . De Lange, Fp. (2013). Human Decision Making Based on Variations in Internal Noise: An EEG Study. *PLoS ONE*, 8 (7) , Article e68928.

Andersen, R., Andersen, K., Hwang, E., & Hauschild, M. (2014). Optic Ataxia: From Balint's Syndrome to the Parietal Reach Region. *Neuron, 81*(5), 967-983.

Annis, D. (1974). Belief and knowledge. *The Philosophical Quarterly, 24*(94), 81-82.

Ansari, D. (2007). Does the Parietal Cortex Distinguish between "10," "Ten," and Ten Dots? *Neuron, 53*(2), 165-167.

Ansari, D., & Vogel, S. (2013). Cognitive Neuroscience of Numerical Cognition. *The Oxford Handbook of Cognitive Neuroscience, Volume 2: The Cutting Edges.* Oxford.

Apker, G., Darling, T., & Buneo, C. (2010). Interacting noise sources shape patterns of arm movement variability in three-dimensional space. *Journal of Neurophysiology, 104(5)*, 2654-66.

Apker, G., & Buneo, C. (2012). Contribution of execution noise to arm movement variability in three-dimensional space. *Journal of Neurophysiology, 107*(1), 90-102.

Arazi, A., Censor, N., & Dinstein, I. (2017). Neural Variability Quenching Predicts Individual Perceptual Abilities. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience, 37*(1), 97-109.

Armstrong, D. (1973). *Belief, truth and knowledge*. Cambridge: University Press.

Augusto, L. (2011). Unconscious knowledge: A survey. *Advances in Cognitive Psychology, 6*, 116-141.

Banks, W., Fujii, M., & Kayra-Stuart, F. (1976). Semantic congruity effects in comparative judgments of magnitudes of digits. *Journal of Experimental Psychology: Human Perception and Performance, 2*(3), 435-447.

Barth, H. & Paladino, A. (2011). The Development of Numerical Estimation: Evidence against a Representational Shift. *Developmental Science, 14*(1), 125-135.

Battaglia-Mayer, A., Caminiti, R., Lacquaniti, F., & Zago, M. (2003). Multiple Levels of Representation of Reaching in the Parieto-frontal Network. *Cerebral Cortex, 13*(10), 1009-1022.

Baud-Bovy, G., & Viviani, P. (1998). Pointing to kinematic targets in space. J Neurosci, 18, 1528–1545.

Beck, J. (2015). Analogue Magnitude Representations: A Philosophical Introduction. *The British Journal for the Philosophy of Science, 66*(4), 829-855.

Behar, I., & Bevan, W. (1961). The perceived duration of auditory and visual intervals: Cross-modal comparison and interaction. *The American Journal of Psychology, 74*, 17-26.

Beilock, S. & Carr, T. (2001). On the Fragility of Skilled Performance: What Governs Choking Under Pressure? *Journal of Experimental Psychology: General, 130*(4), 701-725.

Beran, M., Rumbaugh, J., & Savage-Rumbaugh, D. (1998). Chimpanzee (Pan Troglodytes) Counting in a Computerized Testing Paradigm. The Psychological Record, 48(1), 3-19.

Beran, M., Decker, S., Schwartz, A., & Schultz, N. (2011). Monkeys (Macaca Mulatta and Cebus Apella) and Human Adults and Children (Homo Sapiens) Compare Subsets of Moving Stimuli Based on Numerosity. *Frontiers in Psychology, 2*, Frontiers in Psychology, 2011, Vol.2.

Berteletti, I., Lucangeli, D., Piazza, M., Dehaene, S., Zorzi, M., & García Coll, Cynthia. (2010). Numerical Estimation in Preschoolers. *Developmental Psychology, 46*(2), 545-551.

Boccia, M., Nemmi, F., & Guariglia, C. (2014). Neuropsychology of Environmental Navigation in Humans: Review and Meta-Analysis of fMRI Studies in Healthy Participants. *Neuropsychology Review,24*(2), 236-251.

Bocianski, D., Müsseler, J., & Erlhagen., M (2010). Effects of attention on a relative mislocalization with successively presented stimuli. *Vision Research, 50*(18), 1793-1802.

Bock, O. (1986). Contribution of retinal versus extraretinal signals towards visual localization in goal-directed movements. *Experimental Brain Research, 64*(3), 476-482.

Bock, O. & Daunicht, W. (1987). Information processing in goal-directed movements. *Behavioural Brain Research*, 23(1), 23-28.

Boone, W., & Piccinini, G. (2016). The cognitive neuroscience revolution. *Synthese, 193*(5), 1509-1534.

Bonato, M., Zorzi, M., & Umiltà, C. (2012). When time is space: Evidence for a mental time line. *Neuroscience and Biobehavioral Reviews,36*(10), 2257-2273.

Brannon, E. (2006). The representation of numerical magnitude. *Current Opinion in Neurobiology, 16*(2), 222-229.

Brannon, E., Lutz, D., & Cordes, S. (2006). The development of area discrimination and its implications for number representation in infancy. *Developmental Science, 9*(6), F59-F64.

Brannon, E., Suanda, S., & Libertus, K. (2007). Temporal Discrimination Increases in Precision over Development and Parallels the Development of Numerosity Discrimination. *Developmental Science, 10*(6), 770-777.

Brenner, E., Van Beers, R., Rotman, G., & Smeets, J. (2006). The Role of Uncertainty in the Systematic Spatial Mislocalization of Moving Objects. *Journal of Experimental Psychology: Human Perception and Performance,32*(4), 811-825.

Brenner, E. B., Smeets, J., & Mamassian, P. (2008). If I saw it, it probably wasn't far from where I was looking. Journal of Vision, 8(2), 7.1-10.

Brown, J. (2013). Knowing-how: Linguistics and cognitive science. *Analysis, 73*(2), 220-227.

Brunamonti, E., Falcone, R., Genovesio, A., Costa, S., & Ferraina, S. (2012). Gaze orientation interferes with mental numerical representation. Cognitive Processing, 13(4), 375-379.

Buckley, P. & Gillman, C. (1974). Comparisons of digits and dot patterns. *Journal of Experimental Psychology,103*(6), 1131-1136.

Buckwalter, W. (2014a). FACTIVE VERBS AND PROTAGONIST PROJECTION. *Episteme-A Journal Of Individual And Social Epistemology, 11*(4), 391-409.

Buckwalter, W. (2014b). The Mystery of Stakes and Error in Ascriber Intuitions. In J. Beebe (ed.), *Advances in Experimental Epistemology*, Bloomsbury: 145–74.

Buckwalter, W., and Schaffer, J. (2015). Knowledge, Stakes, and Mistakes. *Noûs*, 49 (2): 201–34.

Bueti, D., & Walsh, V. (2009). The parietal cortex and the representation of time, space, number and other magnitudes. *Philosophical Transactions Of The Royal Society B-Biological Sciences, 364*(1525), 1831-1840.

Butterfill, S., & Sinigaglia, C. (2014). Intention and Motor Representation in Purposive Action. *Philosophy and Phenomenological Research, 88*(1), 119-145.

Cameron, J., director (1997). *Titanic*.

Caminiti, R., Urbano, P.B., Johnson, A., & Caminiti, A. (1990). Making arm movements within different parts of space: Dynamic aspects in the primate motor cortex. Journal of Neuroscience,10(7), 2039-2058.

Cantlon, J., Brannon, E., & Mackintosh, Nicholas. (2007). How Much Does Number Matter to a Monkey (Macaca mulatta) ? *Journal of Experimental Psychology: Animal Behavior Processes,33*(1), 32-41.

Caramazza, A. (1992). Is Cognitive Neuropsychology Possible? *Journal of Cognitive Neuroscience, 4*(1), 80-95.

Carey, S. (2009). *The origin of concepts* (Oxford series in cognitive development). New York; Oxford: Oxford University Press.

Carnap, R. (1928). Der logische Aufbau der Welt. Leipzig: Felix Meiner Verlag.

Carpenter, R. (1990). Neurophysiology (Second ed., Physiological principles in medicine). London: Edward Arnold.

Carrozzo, M., McIntyre, J., Zago, M., & Lacquaniti, F. (1999). Viewer-centered and body-centered frames of reference in direct visuo-motor transformations. Experimental Brain Res, 129, 201–210.

Carter, J., Jarvis, A., & Rubin, B. (2015). Varieties of cognitive achievement. *Philosophical Studies, 172*(6), 1603-1623.

Carter, J., & Pritchard, A. (2016). Perceptual knowledge and relevant alternatives. *Philosophical Studies, 173*(4), 969-990.

Cavina-Pratesi, C., Goodale, M., Culham, J., & He, S. (2007). FMRI Reveals a Dissociation between Grasping and Perceiving the Size of Real 3D Objects (Grasping vs Perceiving Objects). *PLoS ONE,2*(5), E424.

Chang, A., Tzeng, O., Hung, D., & Wu, D. (2011). Big Time Is Not Always Long. *Psychological Science, 22*(12), 1567-1573.

Chen, Y., Monaco, S., Byrne, P., Yan, X., Henriques, D., & Crawford, J. (2014). Allocentric versus egocentric representation of remembered reach targets in human cortex. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience, 34*(37), 12515-26.

Chochon, F., Cohen, L., Moortele, P. F. van de, & Dehaene, S. (1999). Differential Contributions of the Left and Right Inferior Parietal Lobules to Number Processing. *Journal of Cognitive Neuroscience, 11*(6), 617-630.

Churchland, M., Afshar, A., & Shenoy, K. (2006a). A Central Source of Movement Variability. Neuron, 52(6), 1085-1096.

Churchland, M., Yu, B., Ryu, S., Santhanam, G., & Shenoy, K. (2006b). Neural variability in premotor cortex provides a signature of motor preparation. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience,26*(14), 3697-712.

Clark, A. (2001). Visual Experience and Motor Action: Are the Bonds Too Tight? *Philosophical Review, 110*(4), 495-519.

Clark, A. (2007). What Reaching Teaches: Consciousness, Control, and the Inner Zombie. The British Journal for the Philosophy of Science, 58(3), 563-594.

Clark, A. (2009). Perception, action, and experience: Unraveling the golden braid. Neuropsychologia., 47(6), 1460-1468.

Cohen, D., & Blanc-Goldhammer, J. (2011). Numerical bias in bounded and unbounded number line tasks. *Psychonomic Bulletin & Review, 18*(2), 331-338.

Cohen, N., & Squire, L. (1980). Preserved Learning and Retention of Pattern-Analyzing Skill in Amnesia: Dissociation of Knowing How and Knowing that. *Science, 210*(4466), 207-210.

Cohen, N., Eichenbaum, H., Deacedo, B., & Corkin, S. (1985). Different memory systems underlying acquisition of procedural and declarative knowledge. *Annals of the New York Academy of Sciences, 444*, 54-71.

Cohen, S. (1988). How to be a Fallibilist. *Philosophical Perspectives, 2*, 91-123.

Cohen, S. (2002). Basic Knowledge and the Problem of Easy Knowledge. *Philosophy and Phenomenological Research*, 65: 302–29.

Colby, C., & Goldberg, M. (1999). SPACE AND ATTENTION IN PARIETAL CORTEX. *Annual Review of Neuroscience, 22*(1), 319-349.

Coltheart, M. (2001). Assumptions and Methods in Cognitive Neuropsychology. In *The Handbook of Cognitive Neuropsychology*. (eds.) Rapp, B. Philadelphia: Psychology Press.

Cowey, A. (2004). The 30th Sir Frederick Bartlett lecture: Fact, artefact, and myth about blindsight. *The Quarterly Journal of Experimental Psychology Section A, 57*(4), 577-609.

Cressman, E., & Henriques, D. (2010). Reach adaptation and proprioceptive recalibration following exposure to misaligned sensory input. *Journal of Neurophysiology, 103*(4), 1888-95.

Culham, J., Danckert, C., Souza, S., Gati, L., Menon, J., & Goodale, F. (2003). Visually guided grasping produces fMRI activation in dorsal but not ventral stream brain areas. *Experimental Brain Research, 153*(2), 180-189.

Darainy, M., Vahdat, S., & Ostry, D. (2013). Perceptual learning in sensorimotor adaptation. *Journal of Neurophysiology, 110*(9), 2152-62.

Daum, M., Gampe, A., Wronski, C., & Attig, M. (2016). Effects of movement distance, duration, velocity, and type on action prediction in 12-month-olds. *Infant Behavior and Development, 43*, 75-84.

Davidson, D. (1988). Epistemology and Truth, in *Subjective, Intersubjective, Objective*. Oxford: Clarendon Press.

Davis, C., & Kim, J. (2011). What's in a mask? Information masking with forward and backward visual masks. *The Quarterly Journal of Experimental Psychology, 64*(10), 1990-2002.

Davis, H., & Pérusse, R. (1988). Numerical competence in animals: Definitional issues, current evidence, and a new research agenda. *Behavioral and Brain Sciences, 11*(4), 561-579.

DeCharms, R., & Zador, A. (2000). Neural Representation and the Cortical Code. *Annual Review of Neuroscience, 23*(1), 613-647.

De Gelder, B., Tamietto, M., Van Boxtel, G., Goebel, R., Sahraie, A., Van Den Stock, J., . . . Pegna, A. (2008). Intact navigation skills after bilateral loss of striate cortex. *Current Biology: CB, 18*(24), R1128-9.

Dehaene, S., & Akhavein, R. (1995). Attention, Automaticity, and Levels of Representation in Number Processing. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 21*(2), 314-326.

Dehaene, S., & Changeux, J. (1993). DEVELOPMENT OF ELEMENTARY NUMERICAL ABILITIES - A NEURONAL MODEL. *Journal Of Cognitive Neuroscience, 5*(4), 390-407.

Dehaene, S., Spelke, E., Pinel, P., Stanescu, R., & Tsivkin. (1999). Sources of mathematical thinking: Behavioral and brain-imaging evidence. *Science., 284*(5416), 970-974.

Desmurget, M., Vindras, P., Gréa, H., Viviani, P., & Grafton, S. (2000). Proprioception does not quickly drift during visual occlusion. *Experimental Brain Research, 134*(3), 363-377.

Dhawale, A., Smith, M., & Lveczky, B. (2017). The Role of Variability in Motor Learning. *Annual Review of Neuroscience, 40*, 479-498.

Di Bono, M., Casarotti, M., Priftis, K., Gava, L., Umilta, C., & Zorzi, M. (2012). Priming the Mental Time Line. *Journal of Experimental Psychology: Human Perception and Performance, 38*(4), 838-842.

Dinstein, I., Heeger, D., & Behrmann, M. (2015). Neural variability: Friend or foe? *Trends in Cognitive Sciences, 19*(6), 322-328.

Dosher, B., & Lu, Z. (2017). Visual Perceptual Learning and Models. *Annual Review of Vision Science, 3*, 343-363.

Dretske, F. (1970). Epistemic Operators. *The Journal of Philosophy, 67*(24), 1007-1023.

Drew, T., Võ, M., & Wolfe, J. (2013). The Invisible Gorilla Strikes Again. Psychological Science, 24(9), 1848-1853.

Driver, J., Baylis, G., Goodrich, S., & Rafal, R. (1994). Axis-based neglect of visual shapes. *Neuropsychologia, 32*(11), 1353-1356.

Eagleman, D. M., & Sejnowski, T. J. (2007). Motion signals bias localization judgments: A unified explanation for the flash-lag, flash-drag, flash-jump, and Frohlich illusions. *Journal of Vision,7*(4), 3.

Eger, E., Michel, V., Thirion, B., Amadon, A., Dehaene, S., & Kleinschmidt, A. (2009). Deciphering Cortical Number Coding from Human Brain Activity Patterns. *Current Biology, 19*(19), 1608-1615.

Égré, P. (2017). Vague judgment: A probabilistic account. *Synthese, 194*(10), 3837-3865.

Elgin, C. (2004). True Enough. *Philosophical Issues, 14*, 113-131.

Elgin, C. (2017). *True Enough*. Cambridge, MA: MIT Press.

Faisal, A., Selen, A., & Wolpert, L. (2008). Noise in the nervous system. *Nature Reviews., 9*(4), 292-303.

Faivre, N., Berthet, V., & Kouider, S. (2014). Sustained invisibility through crowding and continuous flash suppression: A comparative review. *Frontiers In Psychology, 5*(MAY), 475.

Farrer, C., Franck, N., Paillard, J., & Jeannerod, M. (2003). The role of proprioception in action recognition. *Consciousness and Cognition, 12*(4), 609-619.

Feigenson, Dehaene, & Spelke. (2004). Core systems of number. *Trends in Cognitive Sciences, 8*(7), 307-314.

Ferrari-Toniolo, Papazachariadis, Visco-Comandini, Salvati, D'elia, Di Berardino, . . . Battaglia-Mayer. (2014). A visuomotor disorder in the absence of movement: Does Optic Ataxia generalize to learned isometric hand action? *Neuropsychologia, 63*, 59-71.

Ferretti, G. (2016). Through the forest of motor representations. Consciousness and Cognition, 43, 177-196.

Feigenson, L., Dehaene, S., & Spelke, E. (2004). Core systems of number. *Trends in Cognitive Sciences, 8*(7), 307-314.

Fiehler, K., Rösler, F., & Henriques, D. (2010). Interaction between gaze and visual and proprioceptive position judgements. *Experimental Brain Research, 203*(3), 485-498.

Filimon, F. (2015). Are All Spatial Reference Frames Egocentric? Reinterpreting Evidence for Allocentric, Object-Centered, or World-Centered Reference Frames. *Frontiers in Human Neuroscience., 9*, Frontiers in human neuroscience. , 2015, Vol.9.

Fischhoff, B., Slovic, P., Lichtenstein, S., & Posner, Michael I. (1977). Knowing with certainty: The appropriateness of extreme confidence. *Journal of Experimental Psychology: Human Perception and Performance, 3*(4), 552-564.

Fitch, F. (1963). A Logical Analysis of Some Value Concepts. *J. Symbolic Logic, 28*(Iss.2), 135-142.

Fitelson, B. (2010). 'Strengthening the Case for Knowledge from Falsehood.' Analysis, 70: 666–9.

Flanders, M., Tillery, S., & Soechting, J. (1992). EARLY STAGES IN A SENSORIMOTOR TRANSFORMATION. *Behavioral And Brain Sciences, 15*(2), 309-320.

Fortenbaugh, F., & Robertson, C. (2011). When here becomes there: Attentional distribution modulates foveal bias in peripheral localization. Attention, Perception, & Psychophysics, 73(3), 809-828.

Franz, V., Gegenfurtner, K., Bülthoff, H., & Fahle, M. (2000). Grasping Visual Illusions: No Evidence for a Dissociation Between Perception and Action. *Psychological Science, 11*(1), 20-25.

Franz, V. (2003). Manual size estimation: A neuropsychological measure of perception? *Experimental Brain Research, 151*(4), 471-477.

Franz, V., & Gegenfurtner, K. (2008). Grasping visual illusions: Consistent data and no dissociation. *Cognitive Neuropsychology,25*(7-8), 920-950.

Fraser, L., & Harris, L. (2016). Perceived finger orientation is biased towards functional task spaces. *Experimental Brain Research, 234*(12), 3565.

Fraser, L., & Harris, L. (2017). The effect of hand position on perceived finger orientation in left- and right-handers. *Experimental Brain Research., 235*(12), 3683-3693.

Fridland, E. (2011). The case for proprioception. *Phenomenology and the Cognitive Sciences, 10*(4), 521-540.

Friedman, O. & Turri, J. (2015). Is Probabilistic Evidence a Source of Knowledge? *Cognitive Science, 39*(5), 1062-1080.

Fults, S. (2011). Vagueness and scales. In P. Egré & N. Klinedinst (Eds.), Vagueness and language use (pp. 25–50). New York: Palgrave Macmillan.

Gade, M. Druey, M., Souza, A., & Oberauer, K. (2014). Interference within and between Declarative and Procedural Representations in Working Memory. *Journal of Memory and Language, 76*, 174-194.

Galati, G., Pelle, G., Berthoz, A., & Committeri, G. (2010). Multiple reference frames used by the human brain for spatial perception and memory. *Experimental Brain Research, 206*(2), 109-120.

Gallagher, S. (2003). Bodily self-awareness and object perception. Theoria et Historia Sceintiarum: International Journal for Interdisciplinary Studies, 7(1), 53–68.

Ganel, T. A., Tanzer, M., & Goodale, M. (2008). A double dissociation between action and perception in the context of visual illusions: Opposite effects of real and illusory size: Research report. *Psychological Science, 19*(3), 221-225.

Gallistel, C. & Gelman, R. (2000). Non-verbal numerical cognition: From reals to integers. *Trends in Cognitive Sciences, 4*(2), 59-65.

Gazes, R., Diamond, R., Hope, J., Caillaud, D., Stoinski, T., & Hampton, R. (2017). Spatial representation of magnitude in gorillas and orangutans. *Cognition, 168*, 312-319.

Geary, D., Hoard, M., Nugent, L., & Byrd-Craven, J. (2008). Development of Number Line Representations in Children With Mathematical Learning Disability. *Developmental Neuropsychology, 33*(3), 277-299.

Gemel, A., & Quinon, P. (2015). The approximate numbers system and the treatment of vagueness in conceptual spaces. In P. Lukowski, A. Gemel, & B. Zukowski (Eds.), Cognition, meaning and action. New York: Columbia University Press.

Gerken, M. (2011). Warrant and action. *Synthese, 178*(3), 529-547.

Gerken, M. (2017). On folk epistemology : How we think and talk about knowledge (First ed.).

Gerken, M., & Beebe, J. (2016). Knowledge in and out of Contrast. *Noûs*, 50 (1): 133–64.

Gettier, E. (1963). Is Justified True Belief Knowledge? *Analysis,23*(6), 121-123.

Gibbon, J., & Estes, William K. (1977). Scalar expectancy theory and Weber's law in animal timing. *Psychological Review, 84*(3), 279-325.

Giordani, A., & Mari, L. (2012). Measurement, Models, and Uncertainty. *Instrumentation and Measurement, IEEE Transactions on, 61*(8), 2144-2152.

Glick, E. (2011). Two Methodologies for Evaluating Intellectualism. *Philosophy and Phenomenological Research,83*(2), 398-434.

Göbel, S., Calabria, M., Farnè, A., & Rossetti, Y. (2006). Parietal rTMS distorts the mental number line: Simulating 'spatial' neglect in healthy subjects. Neuropsychologia, 44(6), 860-868.

Goldman, A. (1967). A Causal Theory of Knowing. *The Journal of Philosophy*, *64*(12), 357-372.

Goldman, A. (1988). Psychology and Philosophical Analysis. *Proceedings of the Aristotelian Society, 89*, 195-209.

Goldman, A. (1992). *Liaisons: Philosophy meets the cognitive and social sciences.* Cambridge, Mass.; London: MIT Press.

Goldman, A. (1993a). Epistemic Folkways and Scientific Epistemology. *Philosophical Issues, 3*, 271-285.

Goldman, A. (1993b). *Philosophical applications of cognitive science* (Focus series (Westview Press)). Boulder: Westview Press.

Goldstone, S., Boardman, W., & Lhamon, W. (1959a). Intersensory comparisons of temporal judgments. *Journal of Experimental Psychology, 57*(4), 243-248.

Goldstone, S., Jernigan, C., Lhamon, W., & Boardman, W. (1959b). Further note on intersensory differences in temporal judgment. *Perceptual and Motor Skills, 9*, 252.

Goldstone, S. & Lhamon, W. (1972). Auditory-Visual Differences in Human Temporal Judgment. *Perceptual and Motor Skills,34*(2), 623-633.

Goodale, M. & Milner, D. (1992). Separate visual pathways for perception and action. *Trends in Neurosciences, 15*(1), 20-25.

Greco, John (2012). A (Different) Virtue Epistemology. *Philosophy and Phenomenological Research* 85 (1):1-26.

Grice, H. (1989). *Studies in the way of words*. Cambridge, Mass. ; London: Harvard University Press.

Haffenden, A., & Goodale, M. (1998). The Effect of Pictorial Illusion on Prehension and Perception. *Journal of Cognitive Neuroscience, 10*(1), 122-136.

Halberda, J., Mazzocco, M., & Feigenson, L. (2008). Individual differences in non-verbal number acuity correlate with maths achievement. *Nature, 455*(7213), 665-668.

Halberda, J., Ly, R., Wilmer, J., Naiman, D., & Germine, L. (2012). Number sense across the lifespan as revealed by a massive Internet-based sample. *Proceedings of the National Academy of Sciences, 109*(28), 11116-11120.

Hannon, M. (2013). 'Knows' entails truth. *Journal of Philosophical Research, 38*, 349-366.

Harris, J., Schwarzkopf, D., Song, C., Bahrami, B., & Rees, G. (2011). Contextual illusions reveal the limit of unconscious visual processing. *Psychological Science, 22*(3), 399-405.

Harvey, B.M., Klein, B.P., Petridou, N., & Dumoulin, S.O. (2013). Topographic representation of numerosity in the human parietal cortex. *Science, 341*(6150), 1123-1126.

Hawthorne, J. (2004). *Knowledge and lotteries*. Oxford: Clarendon.

Hawthorne, J., & Stanley, J. (2008). Knowledge and action. *Journal of Philosophy*, 105(10), 571-590.

Hayashi, M., Valli, A., & Carlson, S. (2013). Numerical quantity affects time estimation in the suprasecond range. *Neuroscience Letters, 543*, 7-11.

Haynes, J. & Rees, G. (2005). Predicting the orientation of invisible stimuli from activity in human primary visual cortex. *Nature Neuroscience, 8*(5), 686-91.

Hazlett, A. (2010). The Myth of Factive Verbs. *Philosophy and Phenomenological Research, 80*(3), 497-522.

Hazlett, A. (2012). Factive Presupposition and the Truth Condition on Knowledge. *Acta Analytica, 27*(4), 461-478.

Hellie, B. (2005). Noise and Perceptual Indiscriminability. *Mind,114*(455), 481-508.

He, S. & MacLeod, D. (2001). Orientation-selective adaptation and tilt after-effect from invisible patterns. *Nature, 411*(6836), 473-6.

Helmholtz, H. V. (1866). *Handbuch der physiologischen optik*. Leipzig: Voss.

Hevia, M., Lee, Y., & Streri, A. (2017). The Temporal Dimensions in the First Year of Life. *Timing & Time Perception., 5*(3-4), 280-296.

Higgins, J., & Johnson, M. (2013). Lost Thoughts: Implicit Semantic Interference Impairs Reflective Access to Currently Active Information. *Journal of Experimental Psychology: General, 142*(1), 6-11.

Holton, R. (1997). Some telling examples: A reply to Tsohatzidis. *Journal of Pragmatics, 28*(5), 625-628.

Humphreys, G., Gillebert, C., Chechlacz, M., & Riddoch, M. (2013). Reference frames in visual selection. *Annals of the New York Academy of Sciences, 1296*, 75-87.

Ichikawa, J. & Steup, M. The Analysis of Knowledge. *The Stanford Encyclopedia of Philosophy* (Summer 2018 Edition), Edward N. Zalta (ed.), forthcoming URL = <https://plato.stanford.edu/archives/sum2018/entries/knowledge-analysis/>.

ISO (1993). *Guide to the expression of uncertainty in measurement*. Genève: International Organization for Standardization.

Jacob, S., & Nieder, A. (2009). Tuning to non-symbolic proportions in the human frontoparietal cortex. *European Journal of Neuroscience, 30*(7), 1432-1442.

James, T., Culham, J., Humphrey, G., Milner, A., & Goodale, M. (2003). Ventral occipital lesions impair object recognition but not object-directed grasping: An fMRI study. *Brain, 126*(11), 2463-2475.

Jiang, Y., Costello, P., Fang, F., Huang, M., & He, S. (2006). A gender- and sexual orientation-dependent spatial attentional effect of invisible images. *Proceedings of the National Academy of Sciences, 103*(45), 17048.

Jones, S., Cressman, A., & Henriques, H. (2010). Proprioceptive localization of the left and right hands. *Experimental Brain Research, 204*(3), 373-383.

Jordan, K., MacLean, E., & Brannon, E. (2008). Monkeys Match and Tally Quantities across Senses. *Cognition, 108*(3), 617-625.

Kaplan, M. (2006). If You Know You Can't Be Wrong, in *Epistemology futures*. Oxford: Clarendon.

Karnath, H., Rüter, J., Mandler, A., & Himmelbach, M. (2009). The anatomy of object recognition--visual form agnosia caused by medial occipitotemporal stroke. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience, 29*(18), 5854-62.

Karni, A. (1996). The acquisition of perceptual and motor skills: A memory system in the adult human cortex. *Cognitive Brain Research, 5*(1), 39-48.

Karolis, V., Iuculano, T., & Butterworth, B. (2011). Mapping Numerical Magnitudes along the Right Lines: Differentiating between Scale and Bias. *Journal of Experimental Psychology: General, 140*(4), 693-706.

Kaunitz, L., Fracasso, A., Skujevskis, M., & Melcher, D. (2014). Waves of visibility: Probing the depth of inter-ocular suppression with transient and sustained targets. *Frontiers in Psychology, 5*, 804.

Kelly, T. (2014). Quine and Epistemology, in G. Harman and E. Lepore (eds.), *A Companion to W.V.O. Quine*, Wiley-Blackwell, pp. 17–37.

Kentridge, R. (2015). What is it like to have type-2 blindsight? Drawing inferences from residual function in type-1 blindsight. *Consciousness and Cognition, 32*, 41-4

Kerzel, D. (2002). Memory for the position of stationary objects: Disentangling foveal bias and memory averaging. Vision Research, 42(2), 159-167.

Körding, K., & Wolpert, D. (2004). Bayesian integration in sensorimotor learning. *Nature, 427*(6971), 244-7.

Koivisto, M., & Revonsuo, A. (2007). How Meaning Shapes Seeing. Psychological Science, 18(10), 845-849.

Kim, D., & Opfer, J. (2017). A Unified Framework for Bounded and Unbounded Numerical Estimation. *Developmental Psychology,53*(6), 1088-1097.

Kim, J. (1988). What is Naturalized Epistemology?, in James E. Tomberlin (ed.), *Philosophical Perspectives*, 2, Atascadero, CA: Ridgeview Publishing Co., pp. 381–406.

Knops & Willmes. (2014). Numerical ordering and symbolic arithmetic share frontal and parietal circuits in the right hemisphere. *NeuroImage, 84*, 786-795.

Kornblith, H. (2002). *Knowledge and its place in nature*. Oxford: Clarendon.

Kozuch, B. (2015). Dislocation, Not Dissociation: The Neuroanatomical Argument Against Visual Experience Driving Motor Action. *Mind & Language, 30*(5), 572-602.

Kroustallis, B. (2005) Blindsight, Philosophical Psychology, 18:1, 31-43, DOI: 10.1080/09515080500085346

Lackey, J. (2007). Why we don't deserve credit for everything we know. *Synthese, 158*(3), 345-361.

Lackey, J. (2009). Knowledge and credit. *Philosophical Studies,142*(1), 27-42.

Laeng, B., Chabris, C. F., & Kosslyn, S. M. (2003). Asymmetries in encoding spatial relations. In K. Hugdahl and R. Davidson (Eds.), *The asymmetrical brain* (pp. 303–339). Cambridge, MA: MIT Press.

Laeng, B. (2014). Representation of Spatial Relations. In Ochsner, K., & Kosslyn, S. (Eds.), *The Oxford handbook of cognitive neuroscience* (Oxford library of psychology).

Landy, D., Silbert, N., & Goldin, A. (2013). Estimating Large Numbers. Cognitive Science, 37(5), 775-799.

Landy, D., Charlesworth, A., & Ottmar, E. (2017). Categories of Large Numbers in Line Estimation. *Cognitive Science, 41*(2), 326-353.

Lathrop, W., Bridgeman, B., & Tseng, P. (2011). Perception in the Absence of Attention: Perceptual Processing in the Roelofs Effect during Inattentional Blindness. Perception, 40(9), 1104-1119.

Leibowitz, H., Myers, N., Grant, D., & Melton, Arthur W. (1955). Frequency of seeing and radial localization of single and multiple visual stimuli. *Journal of Experimental Psychology, 50*(6), 369-373.

Leone, M., Salles, A., Pulver, A., Golombek, D., & Sigman, M. (2018). Time drawings: Spatial representation of temporal concepts. *Consciousness and Cognition, 59*, 10-25.

Levy, Y. (2013). Intentional action first. *Australasian Journal of Philosophy, 91*(4), 705-718.

Lewis, D. (1996). Elusive knowledge. *Australasian Journal Of Philosophy, 74*(4), 549-567.

Li, D., Karnath, H., & Rorden, C. (2014). Egocentric representations of space co-exist with allocentric representations: Evidence from spatial neglect. *Cortex; a Journal Devoted to the Study of the Nervous System and Behavior, 58*, 161-9.

Lipton, J., & Spelke, E. (2003). Origins of number sense: Large-number discrimination in human infants.(Author Abstract). *Psychological Science, 14*(5), 396-401.

Littlejohn, C. (2015). Knowledge and Awareness. *Analysis, 75*(4), 596-603.

Loftus, A., Nicholls, M., Mattingley, J., & Bradshaw, J. (2008). Left to Right: Representational Biases for Numbers and the Effect of Visuomotor Adaptation. *Cognition,107*(3), 1048-1058.

Longo, M. & Lourenco, S. (2007). Spatial attention and the mental number line: Evidence for characteristic biases and compression. *Neuropsychologia, 45*(7), 1400-1407.

Lourenco, S., & Longo, M. (2010). General magnitude representation in human infants. *Psychological Science, 21*(6), 873-81.

Luper, S. "Epistemic Closure", *The Stanford Encyclopedia of Philosophy* (Spring 2016 Edition), Edward N. Zalta (ed.), URL = <https://plato.stanford.edu/archives/spr2016/entries/closure-epistemic/>.

Luzzi, F. (2010). 'Counter-Closure.' Australasian Journal of Philosophy, 88: 673–83.

Lyons, I. & Beilock, S. (2011). Numerical ordering ability mediates the relation between number-sense and arithmetic competence. *Cognition, 121*(2), 256-261.

Mack, A. & Rock, I. (1998). *Inattentional blindness* (MIT Press/Bradford books series in cognitive psychology). Cambridge, Mass. ; London: MIT.

Mari, L. (2003). Epistemology of measurement. *Measurement,34*(1), 17-30.

Marin, D., Pitteri, M., Della Puppa, A., Meneghello, F., Biasutti, E., Priftis, K., & Vallesi, A. (2016). Mental Time Line Distortion in Right-Brain-Damaged Patients: Evidence From a Dynamic Spatiotemporal Task. *Neuropsychology*, 30(3), 338-345.

Marr, D. (1982). *Vision: A computational investigation into the human representation and processing of visual information*. New York: W.H. Freeman.

Matthen, M., & Prinz, J. (2015). Unconscious Perception. In *The Oxford Handbook of Philosophy of Perception* (p. The Oxford Handbook of Philosophy of Perception, Chapter 033). Oxford University Press.

McGuire, L. & Sabes, P. (2009). Sensory transformations and the use of multiple reference frames for reach planning. Nature Neuroscience, 12(8), 1056-61.

McIntosh, R. (2010). Optic Ataxia. in EB Goldstein (ed.), Encyclopedia of Perception. Sage Publications Inc., pp. 706-708.

McIntyre, J., Stratta, F., Lacquaniti, F., & McIntyre, F. (1998). Short-term memory for reaching to visual targets: Psychophysical evidence for body-centered reference frames. *Journal of Neuroscience, 18*(20), 8423-8435.

McIntyre, J., Stratta, F., Droulez, J., & Lacquaniti, F. (2000). Analysis of Pointing Errors Reveals Properties of Data Representations and Coordinate Transformations Within the Central Nervous System. *Neural Computation, 12*(12), 2823-2855.

McKendry, Y., Ownsworth, T., & Bettens, G. (2014). Comparing accuracy of knowledge of functional effects of schizophrenia and brain injury. *Psychiatry Research, 219*(1), 225-227.

Milner, A., Perrett, D., Johnston, R., Benson, P., Jordan, T., Heeley, W., Bettucci, D., Mortara, F., Mutani, R., Terazzi, E., & Davidson, D. (1991). PERCEPTION AND ACTION IN 'VISUAL FORM AGNOSIA.' *Brain* 114 (1): 405-428.

Milner, A. & Goodale, M. (1995). The visual brain in action (Oxford psychology series ; no. 27). Oxford: Oxford University Press.

Milner, A. & Goodale, M. (2008). Two visual systems re-viewed. Neuropsychologia, 46(3), 774-785.

Möhring, W., Liu, R., & Libertus, M. (2017). Infants' Speed Discrimination: Effects of Different Ratios and Spatial Orientations. *Infancy., 22*(6), 762-777.

Mohr, H., Mues, H., Robol, V., & Sireteanu, R. (2010). Altered mental number line in amblyopia—Reduced pseudoneglect corresponds to a decreased bias in number estimation. *Neuropsychologia, 48*(6), 1775-1781.

Mole, C. (2009). Illusions, demonstratives, and the zombie action hypothesis. Mind, 118(472), 995-1011.

Montaser-Kouhsari, L., Moradi, F., Zandvakili, A., & Esteky, H. (2004). Orientation-Selective Adaptation during Motion-Induced Blindness. *Perception, 33*(2), 249-254.

Mou, Y. & VanMarle, K. (2014). Two core systems of numerical representation in infants. Developmental Review, 34(1), 1-25.

Moyer, R. & Landauer, T. (1967). Time required for Judgements of Numerical Inequality. *Nature, 215*(5109), 1519-20.

Müsseler, J., Heijden, A., Mahmud, H., Deubel, C., & Ertsey, S. (1999). Relative mislocalization of briefly presented stimuli in the retinal periphery. Perception & Psychophysics, 61(8), 1646-1661.

Mukherjee, T., Battifarano, M., Simoncini, C., & Osborne, L. (2015). Shared sensory estimates for human motion perception and pursuit eye movements. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience, 35*(22), 8515-30.

Nagel, J. (2010). Knowledge Ascriptions and the Psychological Consequences of Thinking about Error. *Philosophical Quarterly*, 60 (239): 286–306.

Nagel, J., San Juan, V., & Mar, R. (2013). Lay Denial of Knowledge for Justified True Beliefs. *Cognition*, 129: 652–61.

Neri, P. (2010). How inherently noisy is human sensory processing? *Psychonomic Bulletin & Review, 17*(6), 802-808.

Nieder, A. (2016). The neuronal code for number. Nature Reviews Neuroscience, 17(6), 366-382.

Oberauer, K. (2010). Declarative and Procedural Working Memory: Common Principles, Common Capacity Limits? *Psychologica Belgica, 50*(3), 277-308.

Ochsner, K., & Kosslyn, S. (2013). Introduction to The Oxford Handbook of Cognitive Neuroscience: Cognitive Neuroscience—Where Are We Now? in *The Oxford Handbook of Cognitive Neuroscience* (1st ed.). Oxford University Press.

Oddie, G. "Truthlikeness", *The Stanford Encyclopedia of Philosophy* (Winter 2016 Edition), Edward N. Zalta (ed.), URL = <https://plato.stanford.edu/archives/win2016/entries/truthlikeness/>.

Odegaard, B., Wozny, D., Shams, L., & Faisal, A. (2015). Biases in Visual, Auditory, and Audiovisual Perception of Space (Biases in Visual, Auditory, and Audiovisual Perception of Space). *11*(12), E1004649.

Oliveri, M., Vicario, C., Salerno, S., Koch, G., Turriziani, P., Mangano, R., Chillemi, G., & Caltagirone, C. (2008). Perceiving numbers alters time perception. *Neuroscience Letters, 438*(3), 308-311.

Opfer, J. E., & DeVries, J. M. (2007). Representational change and magnitude estimation: Why young children can make more accurate salary comparisons than adults. *Cognition*, 108, 843– 849.

Osborne, L., Lisberger, S., & Bialek, W. (2005). A sensory source for motor variation. *Nature, 437*(7057), 412-6.

Owens, D. (2003). Does Belief Have an Aim? *Philosophical Studies, 115*(3), 283-305.

Parent, T. (2014). Knowing- Wh and Embedded Questions. Philosophy Compass, 9(2), 81-95.

Park, J., DeWind, N., Woldorff, M., & Brannon, E. (2016). Rapid and Direct Encoding of Numerosity in the Visual Stream. *Cerebral Cortex, 26*(2), 748-763.

Parkes, L., Lund, J., Angelucci, A., Solomon, J., & Morgan, M. (2001). Compulsory averaging of crowded orientation signals in human vision. *Nature Neuroscience, 4*(7), 739-44.

Parkman, J. (1971). Temporal aspects of digit and letter inequality judgments. *Journal of Experimental Psychology, 91*(2), 191-205.

Pasley, B.N., Mayes, L.C. & Schultz, R.T. (2004) Subcortical discrimination of unperceived objects during binocular rivalry. *Neuron* **42**, 163–172

Pasqualotto, Achille, Taya, Shuichiro, & Proulx, Michael J. (2014). Sensory deprivation: Visual experience alters the mental number line. Behavioural Brain Research, 261, 110.

Peeters, D., Verschaffel, L., & Luwel, K. (2017). Benchmark-based strategies in whole number line estimation. *British Journal of Psychology, 108*(4), 668-686.

Pepperberg, I., & Carey, S. (2012). Grey Parrot Number Acquisition: The Inference of Cardinal Value from Ordinal Position on the Numeral List. *Cognition, 125*(2), 219-232.

Perenin, M., & Vighetto, A. (1988). Optic ataxia: A specific disruption in visuomotor mechanisms. I. Different aspects of the deficit in reaching for objects. *Brain : A Journal of Neurology, 111 (Pt 3)*, 643-74.

Piazza, M., Mechelli, A., Butterworth, B., & Price, C. (2002). Are Subitizing and Counting Implemented as Separate or Functionally Overlapping Processes? *NeuroImage*, *15*(2), 435-446.

Piazza, M., Izard, V., Pinel, P., Le Bihan, D., & Dehaene, S. (2004). Tuning Curves for Approximate Numerosity in the Human Intraparietal Sulcus. *Neuron, 44*(3), 547-555.

Piazza, M., Pinel, P., Le Bihan, D., & Dehaene, S. (2007). A Magnitude Code Common to Numerosities and Number Symbols in Human Intraparietal Cortex. *Neuron, 53*(2), 293-305.

Piazza, M., Facoetti, A., Trussardi, A., Berteletti, I., Conte, S., Lucangeli, D., Dehaene, S. & Zorzi, M. (2010). Developmental Trajectory of Number Acuity Reveals a Severe Impairment in Developmental Dyscalculia. Cognition, 116(1), 33-41.

Piazza, M. & Eger, E. (2016). Neural foundations and functional specificity of number representations. *Neuropsychologia*, *83*, 257-273.

Pinillos, Á. (2012). Knowledge, Experiments, and Practical Interests. In J. Brown and M. Gerken (eds.), *Knowledge Ascriptions*. Oxford: Oxford University Press: 192–220.

Pitts, M., Martínez, A., & Hillyard, S. (2012). Visual processing of contour patterns under conditions of inattentional blindness. *Journal of Cognitive Neuroscience*, *24*(2), 287-303.

Plato. (1921). Plato in Twelve Volumes, Vol. 12 translated by Harold N. Fowler. Cambridge: Harvard University Press; London: William Heinemann Ltd.

Plooy, A., Tresilian, J., Mon-Williams, M., & Wann, J. (1998). The contribution of vision and proprioception to judgements of finger proximity. *Experimental Brain Research, 118*(3), 415-420.

Pöppel, E. (1973). Residual visual function after brain wounds involving the central visual pathways in man. Nature., 243, 295.

Popper, K. (1963). *Conjectures and refutations : The growth of scientific knowledge*. London: Routledge & K. Paul.

Prablanc, C., Echallier, J., Komilis, F., & Jeannerod, E. (1979). Optimal response of eye and hand motor systems in pointing at a visual target. *Biological Cybernetics, 35*(2), 113-124.

Prather, R. (2014). Numerical discrimination is mediated by neural coding variation. *Cognition*, 133(3), 601-610.

Pritchard, D. (2010). Knowledge and Final Value. In *The nature and value of knowledge: Three investigations*. Pritchard, D., Haddock, A., Millar, D., & Millar, Alan. Oxford: Oxford University Press.

Pritchard, D., & Kelp, C. (2010). Two Deflationary Approaches to Fitch-Style Reasoning. In *New Essays on the Knowability Paradox*. Oxford University Press.

Pritchard, D. (2012). Anti-luck virtue epistemology. *Journal of Philosophy, 109*(3), 247-279.

Pritchett, L., Carnevale, M., & Harris, M. (2012). Reference frames for coding touch location depend on the task. *Experimental Brain Research, 222*(4), 437-445.

Quine, W. (1969). Epistemology Naturalized. In *Ontological relativity and other essays* (John Dewey essays in philosophy; no. 1). New York: Columbia University Press.

Quine, W. (1992). *Pursuit of Truth*, revised edition, Cambridge: Harvard University Press.

Ratcliff, R., Voskuilen, C., & McKoon, G. (2018). Internal and external sources of variability in perceptual decision-making. *The Psychological Review., 125*(1), 33-46.

Rajimehr, R., Montaser-Kouhsari, L., & Afraz, S. (2003). Orientation-Selective Adaptation to Crowded Illusory Lines. *Perception, 32*(10), 1199-1210.

Rajimehr, R. (2004). Unconscious Orientation Processing. *Neuron, 41*(4), 663-673.

Rao, S., Mayer, A., & Harrington, D. (2001). The evolution of brain activation during temporal processing. *Nature Neuroscience, 4*(3), 317-323.

Rapp, B. (2001). *The handbook of cognitive neuropsychology: What deficits reveal about the human mind*.

Rauk, M., & Luuk, A. (1980). Identification and detection of spatial position in one-dimensional patterns. Acta et Commentationes Universitatis Tartuensis, 522, Problems of Cognitive Psychology. Tartu, 143-163.

Retsa, C., Naish, P., Bekinschtein, T., & Bak, T. (2016). Temporal judgments in multi–sensory space. *Neuropsychologia, 88*, 101-112.

Revkin, S., Piazza, M., Izard, V., Cohen, L., & Dehaene, S. (2008). Does Subitizing Reflect Numerical Estimation? *Psychological Science, 19*(6), 607-614.

Roggeman, C., Fias, W., & Verguts, T. (2010). Salience maps in parietal cortex: Imaging and computational modeling. *NeuroImage,52*(3), 1005-1014.

Roggeman, C., Santens, S., Fias, W., & Verguts, T. (2011). Stages of nonsymbolic number processing in occipitoparietal cortex disentangled by fMRI adaptation. *JOURNAL OF NEUROSCIENCE*,JOURNAL OF NEUROSCIENCE, 2011.

Ross, A., Schenk, T., & Hesse, C. (2015). The Effect of Gaze Position on Reaching Movements in an Obstacle Avoidance Task.

Rugani, R., Vallortigara, G., Priftis, K., & Regolin, L. (2015a). Animal cognition. Number-space mapping in the newborn chick resembles humans' mental number line. *Science (New York, N.Y.),347*(6221), 534-6.

Rugani, R., Vallortigara, G., Priftis, K., & Regolin, L. (2015b). ANIMAL COGNITION. Response to Comments on "Number-space mapping in the newborn chick resembles humans' mental number line". *Science (New York, N.Y.), 348*(6242), 1438.

Rysiew, P. "Naturalism in Epistemology", *The Stanford Encyclopedia of Philosophy* (Spring 2017 Edition), Edward N. Zalta (ed.), URL = <https://plato.stanford.edu/archives/spr2017/entries/epistemology-naturalized/>.

Saidi, M., Towhidkhah, F., Lagzi, F., & Gharibzadeh, S. (2012). The effect of proprioceptive training on multisensory perception under visual uncertainty. *Journal of Integrative Neuroscience,11*(04), 401-415.

Sanders, M., Warrington, E., Marshall, J., & Weiskrantz, L. (1974). 'Blindsight': Vision in a field defect. *The Lancet (North American Edition)*, 707.

Santens, S., Roggeman, C., Fias, W., & Verguts, T. (2010). Number Processing Pathways in Human Parietal Cortex. *Cerebral Cortex,20*(1), 77-88.

Sartori, G. & Job, R. (1988). The Oyster with Four Legs: A Neuropsychological Study on the Interaction of Visual and Semantic Information. *Cognitive Neuropsychology*, *5*(1), 105-132.

Schnee, I. (2015). There is no knowledge from falsehood. *Episteme, 12*(1), 53-74.

Schneegans, S., & Schöner, G. (2012). A neural mechanism for coordinate transformation predicts pre-saccadic remapping. *Biological Cybernetics, 106*(2), 89-109.

Schnuerch, R., Kreitz, C., Gibbons, H., & Memmert, D. (2016). Not Quite So Blind: Semantic Processing Despite Inattentional Blindness. *Journal of Experimental Psychology: Human Perception and Performance.*

Seegmiller, J., Watson, J., & Strayer, D. (2011). Individual Differences in Susceptibility to Inattentional Blindness. Journal of Experimental Psychology: Learning, Memory, and Cognition, 37(3), 785-791.

Schaffer, J., & Knobe, J. (2012). Contrastive Knowledge Surveyed. *Noûs*, 46 (4): 675–708.

Sheth, B. & Shimojo, S. (2001). Compression of space in visual memory. Vision Research, 41(3), 329-341.

Shi, Y., & Buneo, C. (2009). Exploring the role of sensor noise in movement variability. *Engineering in Medicine and Biology Society, 2009. EMBC 2009. Annual International Conference of the IEEE, 2009*, 4970-4973.

Shi, Y., & Buneo, C. (2012). Movement variability resulting from different noise sources: A simulation study. *Human Movement Science, 31*(4), 772-790.

Siegler, R., & Opfer, J. (2003). The Development of Numerical Estimation: Evidence for Multiple Representations of Numerical Quantity. *Psychological Science, 14*(3), 237.

Siegler, R. S., & Booth, J. L. (2004). Development of numerical estimation in young children. *Child Development*, 75, 428–444.

Simons, D., & Chabris, C. (1999). Gorillas in Our Midst: Sustained Inattentional Blindness for Dynamic Events. *Perception, 28*(9), 1059-1074.

Skagerlund, K., Karlsson, T., & Traff, U. (2016). Magnitude Processing in the Brain: An fMRI Study of Time, Space, and Numerosity as a Shared Cortical System. *Frontiers in Human Neuroscience, 10*(500), 500.

Sokolowski, H., Fias, W., Bosah Ononye, C., & Ansari, D. (2017a). Are numbers grounded in a general magnitude processing system? A functional neuroimaging meta-analysis. *Neuropsychologia,* Neuropsychologia, 2017.

Sokolowski, H., Fias, W., Mousa, A., & Ansari, D. (2017b). Common and distinct brain regions in both parietal and frontal cortex support symbolic and nonsymbolic number processing in humans : A functional neuroimaging meta-analysis. *NEUROIMAGE* 2017.

Sosa, E. (2009). Knowing full well: The normativity of beliefs as performances. *Philosophical Studies, 142*(1), 5-15.

Souman, J. L., Freeman, T. C., Eikmeier, V. O., & Ernst, M. (2010). Humans do not have direct access to retinal flow during walking. *Journal of Vision, 10*(11), 14.

Spelke, E., & Kinzler, K. (2007). Core Knowledge. Developmental Science, 10(1), 89-96.

Squire, L. (2004). Memory systems of the brain: A brief history and current perspective. *Neurobiology of Learning and Memory,82*(3), 171-177.

Sripada, C., and Stanley, J. (2012). Empirical Tests of Interest-Relative Invariantism. *Episteme*, 9: 3–26.

Stanley, J. & Williamson, T. (2001). Knowing How. *Journal of Philosophy, 98*(8), 411-44.

Steglich-Petersen, A. (2009). Weighing the aim of belief. *Philosophical Studies* 145 (3):395-405.

Steglich-Petersen, A. (2017). Weighing the Aim of Belief Again. *Logos and Episteme* 8 (1):141-145.

Sutton, J. (2007). Without Justification. MIT Press.

Tanaka, H., Worringham, C., & Kerr, G. (2009). Contributions of vision–proprioception interactions to the estimation of time-varying hand and target locations. *Experimental Brain Research,195*(3), 371-382.

Tomonaga, M. (2008). Relative numerosity discrimination by chimpanzees (Pan troglodytes): Evidence for approximate numerical representations. *Animal Cognition, 11*(1), 43-57.

Trick, L. (2008). More than superstition: Differential effects of featural heterogeneity and change on subitizing and counting. *Perception & Psychophysics, 70*(5), 743-760.

Tsohatzidis, S. (2012). How to Forget that "Know" is Factive. *Acta Analytica, 27*(4), 449-459.

Turri, J. (2011). Mythology of the Factive. *Logos & Episteme, 2*(1), 141-150.

Turri, J. (2016). Knowledge, certainty, and assertion. *Philosophical Psychology, 29*(2), 293-299.

Turri, J., Alfano, M., and Greco, J. "Virtue Epistemology", *The Stanford Encyclopedia of Philosophy* (Winter 2017 Edition), Edward N. Zalta (ed.), URL = <https://plato.stanford.edu/archives/win2017/entries/epistemology-virtue/>.

Tuthill, J. & Azim, E. (2018). Proprioception. *Current Biology*, *28*(5), R194-R203.

Vallar, G. (1998). Spatial hemineglect in humans. *Trends in Cognitive Sciences, 2*(3), 87-97.

Van Beers, R., Sittig, A., & Denier van der Gon, J. (1998). The precision of proprioceptive position sense. *Experimental Brain Research, 122*(4), 367-377.

Van Beers, R., Sittig, A., & Gon, J. (1999). Integration of proprioceptive and visual position-information: An experimentally supported model. *Journal of Neurophysiology,81*(3), 1355-64.

Van Beers, R., Wolpert, D., & Haggard, P. (2002). When Feeling Is More Important Than Seeing in Sensorimotor Adaptation. *Current Biology, 12*(10), 834-837.

Van Beers, R., Haggard, P., & Wolpert, D. (2004). The role of execution noise in movement variability. *Journal of Neurophysiology, 91*(2), 1050-63.

Van den Bussche, E., Van den Noortgate, W., Reynvoet, B., & Cooper, Harris. (2009). Mechanisms of Masked Priming: A Meta-Analysis. *Psychological Bulletin, 135*(3), 452-477.

Van der Heijden, A., Van der Geest, J., De Leeuw, F., Krikke, K., & Müsseler, J. (1999). Sources of position-perception error for small isolated targets. *Psychological Research, 62*(1), 20-35.

VanMarle, K., & Wynn, K. (2006). Six-month-old infants use analog magnitudes to represent duration. *Developmental Science, 9*(5), F41-F49.

VanMarle, Kristy, & Wynn, Karen. (2009). Infants' Auditory Enumeration: Evidence for Analog Magnitudes in the Small Number Range. *Cognition, 111*(3), 302-316.

Veilleux, L., & Proteau, L. (2011). Congruent visual and proprioceptive information results in a better encoding of initial hand position. *Experimental Brain Research, 214*(2), 215-224.

Vicario, P., Turriziani, K., Caltagirone, O. (2008). Relativistic Compression and Expansion of Experiential Time in the Left and Right Space (Mental Time Number Line). *PLoS ONE, 3*(3), E1716.

Viguier, A., Clément, G., & Trotter, Y. (2001). Distance Perception within near Visual Space. *Perception, 30*(1), 115-124.

Vilarroya, O. (2017). Neural representation. A survey-based analysis of the notion. *Frontiers in Psychology, 8*, .

Vindras, P., Viviani, P., Desmurget, M., & Prablanc, C. (1998). Pointing errors reflect biases in the perception of the initial hand position. *Journal of Neurophysiology, 79*(6), 3290-3294.

Wade N. & Swanston, M. (1996) A general model for the perception of space and motion. Perception 25: 187–194.

Wallhagen, M. (2007). Consciousness and Action: Does Cognitive Science Support (Mild) Epiphenomenalism? *The British Journal for the Philosophy of Science, 58*(3), 539-561.

Wallis, C. (2008). Consciousness, context, and know-how. *Synthese, 160*(1), 123-153.

Walsh, V. (2003). A theory of magnitude: Common cortical metrics of time, space and quantity. *Trends in Cognitive Sciences,7*(11), 483-488.

Ward, E. & Scholl, J. (2015). Inattentional blindness reflects limitations on perception, not memory: Evidence from repeated failures of awareness. *Psychonomic Bulletin & Review, 22*(3), 722-727.

Warfield, T. (2005). 'Knowledge from Falsehood.' Philosophical Perspectives, 19: 405–16.

Wearden, J., Edwards, H., Fakhri, M., & Percival, A. (1998). Why "sounds are judged longer than lights": Application of a model of the internal clock in humans. *Quarterly Journal Of Experimental Psychology Section B-Comparative And Phys, 51*(2), 97-120.

Weiskrantz, L. (1986). *Blindsight: A case study and implications* (Oxford psychology series; no. 12). Oxford: Clarendon Press.

Westwood, D., & Goodale, M. (2011). Converging evidence for diverging pathways: Neuropsychology and psychophysics tell the same story. *Vision Research, 51*(8), 804-811.

Whiting, D. (2014). Reasons for Belief, Reasons for Action, the Aim of Belief, and the Aim of Action. In Clayton Littlejohn & John Turri (eds.), *Epistemic Norms*. Oxford University Press.

Wilkinson, F. R., Ellemberg, D., & Wilson, H. (1997). Lateral interactions in peripherally viewed texture arrays. *Journal of the Optical Society of America A: Optics and Image Science, and Vision, 14*(9), 2057-2068.

Williams, M., Morris, A., McGlone, F., Abbott, D., & Mattingley, J. (2004). Amygdala Responses to Fearful and Happy Facial Expressions under Conditions of Binocular Suppression. *Journal of Neuroscience, 24*(12), 2898-2904.

Wittgenstein, L. (1953). *Philosophical investigations*. Oxford: B. Blackwell.

Xu, F. & Spelke, E. (2000). Large number discrimination in 6-month-old infants. Cognition, 74(1), B1-B11.

Xu, F., Spelke, E., & Goddard, S. (2005). Number sense in human infants. Developmental Science, 8(1), 88-101.

Xuan, B., Zhang, D., Chen, X., & He, S. (2007). Larger stimuli are judged to last longer. *Journal of Vision, 7*(10), Journal of Vision, 6 July 2007, Vol.7(10).

Yavari, F., Towhidkhah, F., Ahmadi-Pajouh, M., Darainy, M. (2015). The role of internal forward models and proprioception in hand position estimation. *Journal of Integrative Neuroscience,14*(03), 403-18.

Zaehle, T., Jordan, K., Wustenberg, T., Baudewig, J., Dechent, P., Mast, F., & Wuestenberg, T. (2007). The neural basis of the egocentric and allocentric spatial frame of reference. *Brain Research, 1137*(1), 92-103.

Zorzi, M., Di Bono, M., & Fias, W. (2011). Distinct representations of numerical and non-numerical order in the human intraparietal sulcus revealed by multivariate pattern recognition. *NeuroImage,56*(2), 674-680.