University of Southampton

Faculty of Arts and Humanities

Philosophy

A New Theory of Intuition:

How Metacognitive Experiences Shape Epistemology

by

Benjamin Francis Paget-Woods, BA (Hons), BPhil

Thesis for the degree of Doctor of Philosophy

September 2020

University of Southampton

Abstract

Faculty of Arts and Humanities

Philosophy

Thesis for the degree of Doctor of Philosophy

A New Theory of Intuition:

How Metacognitive Experiences Shape Epistemology

by

Benjamin Francis Paget-Woods, BA (Hons), BPhil

It's widely assumed that intuitions are central to the methods of contemporary analytic philosophy. In particular, it's thought that philosophers appeal to intuitions as evidence, or as a source of evidence, for their claims. Indeed, this view, which has become known as 'centrality', has been put forward explicitly by, for example, Chalmers, Kornblith, Bealer, Baz, Richard, and Liao, to name but a few. Recently, however, this interpretation of philosophical practice has been challenged, most notably by Williamson, Deutsch, Ichikawa, and Cappelen (the 'anti-centralists'), who argue that intuitions aren't, after all, central to our arguments. Alongside this debate has come a resurgence of interest in the related question of how philosophers use 'intuition-talk', namely words like 'intuition' and 'intuitive'; if this language isn't citing evidence, then what is its purpose, if anything, and, if it is, then what exactly is being referenced in support of our theories, or, what are intuitions?

In this thesis, I make two, primary claims. First, I argue that intuitions do, indeed, play a central role in analytic philosophy (*contra* the anti-centralists), and help to clarify that role. Specifically, I make the case for a centralist interpretation of the primary argument for epistemic contextualism, identifying, through a conscientious analysis of the most seminal literature, not one but several, specific ways in which contextualists appeal to intuitions in an evidential capacity. Since contextualism is chiefly motivated by said argument, and is having a burgeoning influence on modern epistemology, and considering that epistemology is of ubiquitous philosophical significance, with ties to arguably all other core philosophical topics, such as ethics and metaphysics, I thereby demonstrate that treating intuitions as evidence is profoundly shaping the discipline at large.

Second, I develop a novel account of what the relevant philosophical intuitions are. I argue against extant 'minimalist' theories, showing that they aren't reducible to beliefs or credences of any kind (*contra* Lewis, Parsons, and Kornblith, for example), dispositions to believe (*contra* Sosa and Lycan, for example), temptations to believe (*contra* Williamson and Van Inwagen, for example), or facts about ordinary language. Moreover, I argue that existing accounts of intuitions as 'intellectual seemings' – advocated by, for instance, Brogaard, Huemer, and Bealer – are too conservative to capture the intuitions in question. In place of these alternatives, I propose a more liberal version of the intellectual seemings thesis. Then, I argue that such seemings, and thus the relevant philosophical intuitions, aren't *sui generis*, as many are wont to assume, but are, rather, a sub-category of mental states known in psychology as 'epistemic feelings'. Epistemic feelings are experiences triggered by metacognitive monitoring and control subsystems, in response to features of a first-order cognitive process and/or its outputs, such as its fluency. This interdisciplinary thesis revolutionises our understanding of philosophical intuitions, and bridges two, hitherto largely segregated academic sub-disciplines.

I conclude, overall, that metacognitive experiences profoundly shape philosophy, and briefly consider some of the possible implications of my discovery. In particular, I suggest that we should treat intuitions as higher-, not first-order evidence for their content.

Table of Contents

Table of Figures

Research Thesis: Declaration of Authorship

Print name: Benjamin Francis Paget-Woods

Title of thesis: A New Theory of Intuition: How Metacognitive Experiences Shape Epistemology

I declare that this thesis and the work presented in it are my own and has been generated by me as the result of my own original research.

I confirm that:

1. This work was done wholly or mainly while in candidature for a research degree at this University;
2. Where any part of this thesis has previously been submitted for a degree or any other qualification at this University or any other institution, this has been clearly stated;
3. Where I have consulted the published work of others, this is always clearly attributed;
4. Where I have quoted from the work of others, the source is always given. With the exception of such quotations, this thesis is entirely my own work;
5. I have acknowledged all main sources of help;
6. Where the thesis is based on work done by myself jointly with others, I have made clear exactly what was done by others and what I have contributed myself;
7. None of this work has been published before submission.

Signature: ..... Date: 07.03.2021

Acknowledgements

I could not have completed this thesis without the input and encouragement of others.

In particular, I would like to thank my supervisory team – Dr Conor McHugh, Dr Nat Hansen, and Dr Kurt Sylvan – for their invaluable feedback and advice, and Dr Wesley Buckwalter, for agreeing to be my external examiner. I am also extremely grateful to the Arts and Humanities Research Council and the Royal Institute of Philosophy for funding my research.

In addition, I would like to thank Dr Lee Walters and Dr Kurt Sylvan (again) for their helpful feedback during and after my Third Year Milestone.

I would also like to acknowledge the Philosophy Department at the University of Southampton, its staff, and students, for making me feel so welcome throughout my many years of study here.

Charlotte Unruh, Sigmund Bruno Schilpzand, Dr Alex Gregory, Prof Daniel Whiting, Prof Genia Schönbaumsfeld, and several anonymous reviewers have my gratitude for providing feedback, in various forms, on sections of this work that have been adapted into research papers and presentations.

I would also like to thank my parents, Anne and Michael Paget-Woods, for their unwavering faith in me.

Finally, special thanks go to Carla Wiggs for her endless love, support, and encouragement, and for proofreading the final draft of this thesis.

Definitions and Abbreviations

AI	The argument from ignorance
AIT	The argument from intuition-talk
BIV	Brain-in-a-vat
CPK	The closure principle for knowledge
EC	Epistemic contextualism
E-feelings	Epistemic feelings
FOE	Feeling of error
FOF	Feeling of forgetting
FOK	Feeling of knowing
FOR	Feeling of rightness
IDR	Intuition-driven romanticism
IRI	Interest relative invariantism
JTB	The justified, true belief theory of knowledge
K-sentences	Knowledge-attributing and knowledge-denying sentences
K-words	'Knows' and its cognates
NN	Naming and Necessity
PWI	Philosophy Without Intuitions
RE	Reflective equilibrium
SSI	Subject sensitive invariantism
TOT	The tip-of-the-tongue feeling
X-Phi	Experimental philosophy

Chapter 1 Intuitions and Philosophy

It's widely assumed that intuitions are central to the methods of contemporary analytic philosophy. In particular, it's thought that philosophers appeal to intuitions as evidence, or as a source of evidence, for their claims. Recently, however, this interpretation of philosophical practice has been challenged by the 'anti-centralists', who argue that intuitions aren't, after all, central to our arguments. Alongside this debate has come a resurgence of interest in the related question of how philosophers use 'intuition-talk', namely words like 'intuition' and 'intuitive'; if this language isn't citing evidence, then what is its purpose, if anything, and, if it is, then what exactly is being referenced in support of our theories, or, what are intuitions?

In this thesis, I make two, primary claims. In the first half, I argue that intuitions do, indeed, play a central role in analytic philosophy (*contra* the anti-centralists), and help to clarify that role. Then, in the second half, I develop a novel account of what those intuitions are. I close by briefly considering some of the possible implications of my discoveries.

In this opening chapter, I first outline my research objectives in more detail (§1.1). After that, I explain their motivations, and contextualise them within the extant literature (§1.2). Then, I explain what I take to be the shortfall of the existing research into these matters, and how my own approach will differ (§1.3). I close by outlining my plan for the thesis, offering a brief synopsis of each subsequent chapter (§1.4).

1.1 Research Objectives

To reiterate, it's widely assumed that intuitions play an important role in contemporary analytic philosophy. In particular, it's thought that many philosophers appeal to intuitions as evidence, or as a source of evidence, for their claims. Below, I provide several quotes that demonstrate the popularity of this view.

> "I find it fairly obvious that many philosophers, including myself, appeal to intuitions" (Chalmers, 2014, p.535).

> "George Bealer does it. Roderick Chisholm does it a lot. Most philosophers do it openly and unapologetically, and the rest arguably do it too, although some of them would deny it. What they all do is appeal to intuitions in constructing, shaping, and refining their philosophical views" (Kornblith, 1998, p.129).

> "For several decades now philosophers in the mainstream of analytic philosophy, in their pursuit of a theory of x… have centrally relied on what they themselves have been happy to describe as their own and other people's 'intuitions' of whether or not our concept of x, or 'x', applies to this or that particular case, real or imagined" (Baz, 2012, p.316).

> "Much, though by no means all, philosophy involves the generation of ingenious cases about which philosophers have relatively strong and consistent intuitions; such intuitions are a significant source of evidence for philosophical analysis" (Richard, 2019, p.50).

> "When philosophers run out of arguments, they often appeal to intuitions. If something is intuitive, this tends to count in favor of a position, and if something is counterintuitive, this tends to count against the position" (Liao, 2008, p.248)[1].

One might, for instance, think

- that the evidence against the 'justified true belief' theory of knowledge is meant to be the *intuition* that a subject in a 'Gettier case' truly and justifiably believes, but doesn't know, a certain proposition (Gettier, 1963);
- that the evidence for there being a morally relevant difference between doing and merely allowing harm is meant to be our *intuition* about Foot's 'trolley cases' (Foot, 1967);
- that the evidence for externalism about mental content is meant to be the *intuition* that, on Putnam's 'Twin Earth', 'water' means XYZ, not H_2O (Putnam, 1975);
- that the evidence against materialism is meant to be the *intuition* that physical duplicates of sentient beings which lack phenomenal consciousness (or 'philosophical zombies') are metaphysically possible (Chalmers, 1996);
- and so on.

Recently, this view of philosophical practice has become known, in some circles, as 'centrality'.

[1] In addition, Bealer (1996) calls the method of supporting claims by their conformity to intuition analytic philosophy's "standard justificatory procedure" (p.4).

> *Centrality*: Contemporary analytic philosophers rely on intuitions as evidence (or as a source of evidence) for philosophical theories[2,3].

It's important to understand that, by itself, centrality is a purely descriptive claim, not a normative or evaluative one. To wit, it's a claim to an existing practice within philosophy, not a claim that this practice is one we ought to adopt, nor does it embed any judgement, positive or negative, of that convention. Consequently, holding a view on centrality does not, *ipso facto*, commit one to either an endorsement or condemnation of the modern analytic tradition. Accordingly, whenever I speak of philosophers, say, 'advocating' centrality, I only meant that they argue for the truth of this description, not that they, necessarily, support the use of intuitions in philosophy.

Recently, there's been a veritable explosion of interest in centrality. In particular, philosophers have undertaken to answer four, interrelated metaphilosophical questions:

A) Is centrality actually true?[4]
B) What are intuitions?[5]
C) Are the claims philosophers make about what people intuit accurate?[6]
D) Are intuitions good evidence?[7]

We can roughly divide the above questions into two sets. A and B represent what we might loosely call the 'exegetical project', where the goal is to accurately interpret current philosophical methods: are philosophers relying on intuitions and, if so, what exactly are those intuitions that they're relying on? C and D, meanwhile, represent what we might loosely call the 'evaluative project', where the goal is to assess the merits of those methods: do people really have the alleged intuitions, and, if they do, are those intuitions good evidence for philosophical claims?

Generally, the exegetical project takes precedence over the evaluative one. If philosophers aren't relying evidentially on intuitions, then nothing hinges on whether their claims about what people intuit are accurate, or whether those intuitions would be good evidence for philosophical

---

[2] This formulation of 'centrality' comes from Cappelen (2012, p.77), who popularised the term. Cf. Nado (2015): "philosophers make use of claims about intuitions as premises in their arguments, either explicitly or implicitly" (p.790). Weinberg *et al*. (2001), meanwhile, call the alleged practice 'intuition-driven romanticism', or IDR.

[3] No centralist wants to claim that every philosopher appeals to intuitions in every argument. Rather, we should interpret centrality as saying that appealing to intuitions is a general, but not necessarily universal, practice. Consequently, centrality is compatible with some philosophers never appealing to intuitions.

[4] See, for example, Cappelen (2012), Chalmers (2014), Deutsch (2015), Fischer and Collins (2015), Ichikawa (2014), and Williamson (2011).

[5] See, for example, Brogaard (2013c), Sosa (2007b), Bengson (2014, 2015), and Richard (2014).

[6] See, for example, Stich (2013), Buckwalter (2010, 2012), Buckwalter and Schaffer (2015), Sripada and Stanley (2012). Stanley (2004), Feltz and Zarpentine (2010), and Pinillos (2012).

[7] See, for example, Fischer (2014), Fischer *et al*. (2015), Nichols *et al.* (2003), and Lynch (2006).

theories. Cappelen (2012), for example, who argues that centrality is false, concludes that, consequently, “there’s no urgency in figuring out what intuitions are and what epistemic status they have.” (p.18). Similarly, until we know what ‘intuitions’ – as philosophers are using the term – actually are, we won’t know what to measure when trying to settle the matter of what people ‘intuit’, and we’ll have a hard time making an informed evaluation of their epistemic features.

In this thesis, I address these foremost exegetical affairs. In Chapter 2, as a partial answer to question A, I argue that intuitions do, indeed, play a central role in, at least, some core philosophical arguments, and help to clarify that role. Specifically, I make the case for a centralist interpretation of the primary argument for epistemic contextualism, identifying, through a conscientious analysis of the most seminal literature, not one but several, specific ways in which contextualists appeal to intuitions in an evidential capacity. Moreover, since contextualism is chiefly motivated by said argument, and is having a burgeoning influence on modern epistemology, and considering that epistemology is of ubiquitous philosophical significance, with ties to arguably all other core philosophical topics, such as ethics and metaphysics, I thereby demonstrate that treating said intuitions as evidence is profoundly shaping the discipline at large. To put it another way, arguments that rely on intuitions as evidence have played a substantial role in motivating epistemologists to adopt contextualism, which has had ‘downstream’ effects on large swathes of philosophical thought as a whole, and thus relying on intuitions is and has been measurably shaping the discipline.

I said that this constitutes only a ‘partial’ answer to question A, because the fact that some philosophers have relied evidentially on intuitions, in some arguments, for some theories, which have had a measurable impact on philosophy as a whole, does not entail that philosophers as a whole are relying evidentially on intuitions, nor does it say anything to the effect that they are not, and so doesn’t really qualify as an answer to the matter of whether centrality, which I defined earlier as precisely such a generalisation, is true. However, I do think that the practice I identify in contextualist literature establishes that the *spirit* of centrality is, indeed, correct. To elaborate, centrality is assumed to be an important and interesting truth at least in part because of the metaphilosophical importance that it grants to the study of intuitions and intuition-talk. Put simply, if centrality is true and philosophers in general really are relying on intuitions as a sort of evidence, then it matters what intuitions actually are and if that thing is a reliable guide to philosophical truths: the answers to these questions will have direct and obvious implications for whether current philosophical methods are sound. Indeed, for the most part, the debate over the truth of centrality is largely over whether studies on intuitions are metaphilosophically relevant or not (as we’ll see in the next section). Those who deny the truth of centrality will also tend to deny that studies on intuition have much, if any, metaphilosophical significance (see, for instance,

Cappelen (2012)). However, for such studies to have significance, it isn't necessary for the generalisation which I earlier called 'centrality' to be true. It is enough if intuitions have been, and are by all accounts are continuing to be relied upon in arguments and debates that have had far reaching influence on central domains of philosophy. For, if that is the case – which, to reiterate, I will argue that it is – then, even if relying evidentially on intuitions is *not* a common practice across philosophy as a whole (which, of course, it still might be), whether that methodology is sound is still very important for assessing the state of contemporary philosophical thought. That is, for the study of intuitions to be metaphilosophically significant, it is sufficient for the spirit of centrality, if not the popularised formulation of it that I used earlier, to be true: that intuitions be of 'central' significance to the history and evolution of modern philosophy.

Another way to put this is to say that the formulation of centrality that has become popular in the literature, and which I provided above, doesn't really capture the central issue at stake in the debate. It isn't really a discussion about whether relying on intuitions is a common practice in philosophy, but about whether the study of intuitions is metaphilosophically relevant. Lots of philosophers relying evidentially on intuitions is one way to make it so that it is, but, as I'm arguing here, another way is for relying on intuitions, even *if* it's a minority practice, to be substantially influencing modern philosophical thought.

So, in summation, the first half of the thesis aims to establish that centrality is true in as much as it needs to be to make research questions B, C, and D a worthwhile academic endeavour, and it achieves this, not by making the case for a generalisation, but by showing that the core argument for contextualism, at least, appeals evidentially to something its advocates label 'intuitions'.

Then, in the second half of the thesis, as a partial answer to question B, I develop a novel account of what these intuitions, at least, are. In Chapter 3, I argue against the extant 'minimalist' views that such intuitions are outright beliefs, partial beliefs, dispositions to believe, or facts about ordinary language. I also argue, in Chapter 4, that existing accounts of intuitions as 'intellectual seemings' are too conservative to capture the intuitions in question. In place of these alternatives, I defend a more liberal version of the intellectual seemings thesis. Then, in Chapter 5, I argue that said seemings, and hence certain intuitions, are a sub-category of what some psychologists call 'epistemic feelings'. Epistemic feelings are experiences triggered by metacognitive monitoring and control subsystems, in response to features of a first-order cognitive process and/or its outputs, such as its fluency. This interdisciplinary thesis revolutionises our understanding of (some kinds of) philosophical intuitions, and bridges two, hitherto largely segregated academic sub-disciplines. It also serves as an addendum to the extant taxonomy of philosophical theories of intuition.

In Chapter 6, I conclude, overall, that metacognitive experiences profoundly shape philosophy.

It's important to recognise, therefore, that both halves of this thesis are descriptive, not normative, in nature. Specifically, the first half argues that some philosophers, in some very important and influential arguments, do rely evidentially on intuitions, but does not pass any judgement, positive or negative, on said practice, while the second half argues that the intuitions in question are epistemic feelings, but does not say anything about whether said feelings are reliable, or otherwise a good source of evidence for philosophical theories. Accordingly, the thesis also passes no judgement on contextualism, or on the philosophers and arguments that advocate it by appealing to intuitions. As such, the reader should not proceed with the expectation that such a judgement will arrive. In fact, it cannot be emphasised enough that the plausibility of contextualism is entirely incidental to this research project. The reasons for focussing on this literature specifically are explained in §1.3. For now, it will suffice to see my analysis of contextual arguments as a case study that yields helpful insights; through looking closely at these philosophers, and how they defend their views, I will discover that centrality is true, in spirit, and also uncover many clues about what such philosophers actually mean by 'intuitions'. But the goal is never to pass judgement on whether the methodology so identified is sound. Again, the analysis of contextualism is purely descriptive.

However, the purely descriptive content of this thesis does not undermine its significance. To reiterate, having an accurate description of philosophical practice is necessary for making accurate normative judgements about it. At the very least, it will be much harder for us to fairly assess what philosophers are doing if we only have a false and/or muddled picture in our heads of what that is. In other words, to ever have accurate *normative* judgements about philosophy, we first need accurate *descriptive* judgements about it. In fact, the centrality debate itself serves as a very simple example of why this is. If one were to pass negative judgement on philosophers for using intuitions as evidence, because they think that intuitions are a very poor guide to what's true, prior to having thoroughly researched what intuitions even are, or whether philosophers really do rely on them as evidence, they would risk making egregious errors; on inspection, it may transpire that philosophers don't really rely on intuitions, or that they do, but that said intuitions aren't what the cynic initially assumed they were, but are instead something actually quite reliable. To make normative judgements before arriving at a sufficiently detailed description of the thing being judged is to 'jump the gun', so to speak. To put it simply then, the purpose of this thesis is to help equip metaphilosophers with just such a descriptive picture of what intuitions are and how philosophers actually use them. We can then proceed, in future research, to a more informed normative assessment of philosophical practice. Therefore, while this thesis is descriptive in content, normative implications will follow from these descriptive claims.

Most of the project's normative implications will follow when my claims are taken in combination with other claims, derived from other research projects. For instance, once the descriptive fact that some philosophers rely evidentially on epistemic feelings is established, we can then combine this with other discoveries about the reliability of said feelings to assess the viability of that philosophical method. There simply isn't enough space (by far) to pursue this next stage of the research in this thesis, but my work here lays a solid groundwork for such future endeavours.

All that being said, I *will* briefly speculate about some of the normative implications of my findings in Chapter 6, to give the reader a taste of what future ideas the work done here may yield. For example, I will suggest that we should treat some intuitions as higher-, not first-order evidence for their content. I only mean to caution the reader against interpreting the thesis at large as normative rather than descriptive.

So, in summation, this thesis aims to partially answer research questions A and B. However, I'll first explain, more fully, the motivation behind addressing these two exegetical questions, expanding upon some ideas that I briefly mentioned above, and contextualise my own work within the extant literature (§1.2). Then, I'll explain what I take to be the shortfall in the existing research into these matters, and how my own approach will differ (§1.3).

1.2 Background and Motivation

In this section, I explain the motivation behind answering the two exegetical questions, A (§1.2.1) and B (§1.2.2), and contextualise my own research within the extant literature.

1.2.1 Is Centrality Actually True?

Because centrality has been so widely assumed, and often implicitly, there's been little motivation to provide explicit arguments in its defence. Consequently, such arguments are rare. Indeed, some doubt that they even existed at all prior to the publication of Cappelen's *Philosophy Without Intuitions* (hereafter PWI) in 2012 (although Cappelen ultimately argues against centrality in PWI, he speculates therein about how one might try to defend it). As Cappelen (2015) himself puts it:

> "[N]o *one* has *ever* presented a detailed case for Centrality. I mean this literally: not even a page is devoted to setting out a careful case for a positive answer—it's just assumed that the answer is 'Yes'" (p.577, emphasis in original).

If it's been accepted so unquestioningly until now, it's natural to worry that, on reflection, it might transpire to be false. Indeed, centrality has recently been called into doubt by a small but growing group of influential philosophers, whom I will herein refer to as the 'anti-centralists'. Some anti-

centralists merely claim that, while some philosophers, like Chalmers, explicitly admit to using intuitions to motivate their philosophical theories, the popularity of such approaches has been greatly exaggerated (certainly to a point where a generalisation like centrality is false). Some, however, go further, accusing even philosophers like Chalmers of (unintentionally) misrepresenting their methodology. Deutsch (2010), for example, says:

> "it is common… to misrepresent arguments in philosophy as *depending on* the intuitiveness, or counterintuitiveness, of some proposition or other" (p.448, emphasis in original).

Deutsch goes on to clarify that he's talking about a misrepresentation by philosophers of their own method – reflecting a shortfall in their self-knowledge – and also doubts whether, on reflection, the same individuals would really consider empirical data about the content of people's intuitions to provide a refutation or confirmation of a philosophical thesis.

While this anti-centralism constitutes only a minority view, its proponents include multiple respected figures within the analytic discipline. In addition to Deutsch, the most prominent advocates of this position are the aforementioned Cappelen (2012), who argued at length against centrality in PWI, Ichikawa (2014), who believes that

> "it is at best less than obvious that those psychological states that are intuitions have an important role to play in the epistemology of armchair philosophy" (p.238)

and Williamson (2011), who claims that we don't form judgements about thought experiments using

> "a mysterious *sui generis* faculty of rational intuition, or anything of the kind" but rather "more mundane skills, such as careful attention to details in the description of the scenario and their potential relevance to the questions at issue" (p.217).

This camp finds itself in direct opposition to the many 'centralists', such as the aforementioned Chalmers, Kornblith, Baz, Richard, and Liao, and to the prevailing interpretation of analytic practice that those figures espouse.

Because it's so casually and widely accepted, if centrality does turn out to be false, this discovery could have profound and far-reaching ramifications. I can think of two ways, in particular, in which it would affect present philosophical research.

First, it would require us to reinterpret the seminal arguments for many influential theses in philosophy, which, in turn, will also change how we critically engage with said arguments. In other

words, it will cast doubt on significant amounts of secondary literature, which has assumed a centralist interpretation of those seminal works in its discussion and evaluation of them. Consider, for instance, Gettier's (1963) case against the 'justified true belief' theory of knowledge (JTB). It's commonplace to suppose that Gettier's famous argument against JTB is based on an appeal to intuition, namely the intuition that, in Gettier cases, S truly and justifiably believes, but doesn't know, *p*. At least, that's the version of the argument that's been repeated, and relied upon, time and again in the subsequent literature. Pust (2000), for example, when discussing a classic Gettier thought experiment, says:

> "Most philosophers take the fact that they have the intuition that S does not know that *p* in this case to show that S does not know that *p*" (p.5).

If it transpires that intuition isn't, after all, the evidence provided against JTB, this interpretation has essentially missed the point of the original argument. At the very least, the tendency to read such seminal arguments in a centralist manner might cause contemporary philosophers to overlook the actual historical contribution that their advocates were making, while we engage with a, merely imagined, empirical-psychological premise.

In fact, there may even be reason to think that a centrality reading of Gettier's argument is, indeed, misleading. As Deutsch (2010) notes, Gettier doesn't mention 'intuition' even once in his presentation of the famous thought experiments, or 'Gettier cases'. Rather, he simply states, as fact, that, although Smith has a justified, true belief that (e) (that the man who will get the job has ten coins in his pockets), "it is equally clear that Smith does not know that (e) is true" (Gettier, 1963, p.122). In other words, it's simply assumed, not read, that certain intuitions play an important role in Gettier's argument (or that 'it is equally clear' is meant to be synonymous with 'it is intuitive').

Another, similar example is Putnam's seminal argument for semantic externalism. In his famous 'Twin Earth' argument, it's widely assumed that Putnam is appealing to an intuition: the intuition that, on Twin Earth, 'water' refers to XYZ. Yet, again, we find, on closer inspection, that the allegedly 'intuitive' claim – that, on Twin Earth, "water" refers to XYZ, not to H_2O – is never once called 'intuitive', but simply stated as fact (as Cappelen (2014) has likewise observed):

> "One of the peculiarities of Twin Earth is that the liquid called "water" is not H2O but a different liquid whose chemical formula is very long and complicated. I shall abbreviate this chemical formula simply as XYZ" (Putnam, 1975, p.99).

In other words, that, on Twin Earth, 'water' refers to XYZ is just a fact about that possible world, not an 'intuition' that we're meant to have about it. Again, this suggests that we may have long been missing the real point of Putnam's argument.

To clarify, I'm not arguing here that a centralist reading of these two seminal arguments is mistaken. I'm merely highlighting why the refutation of centrality would be important, and what *prima facie* motivation we have (such as the lack of intuition-talk in these examples) to consider that it might, indeed, be false.

Second, the refutation of centrality would invite serious concerns about large swathes of 'experimental philosophy' (or X-Phi). Many experimental philosophers, or 'experimentalists', gather data about what intuitions people, especially laypeople, actually have, in order to test the accuracy of philosophers' empirical-psychological claims about folk intuitions. Many others focus on gathering higher-order evidence about the epistemic properties of intuitions, usually to the end of evaluating philosophical method. For example, they test for whether intuitions display sensitivity to otiose parameters, such as gender, ethnicity, or socioeconomic background, and argue that, if they do, then they aren't 'truth-tracking', and thus aren't a good source of evidence for philosophical theses. Fischer (2018) summarises the latter approach thusly:

> "This research program uses surveys and experiments to study the sensitivity of intuitions to truth-irrelevant parameters… and infers lack of evidentiary value where it observes such sensitivity" (p.413).

Both projects are, often explicitly, motivated by the belief that intuitions play a central role in the methods of contemporary analytic philosophy (see, for example, Weinberg *et al.*, 2001, p.432). However, if the experimentalists are wrong in this centralist diagnosis, then much of the work they've conducted has been misguided, and is, arguably, incidental: due to the limited role of intuitions in our arguments, questions about what intuitions people have and whether they can be trusted just aren't particularly relevant to first-order analytic practice. At least the majority of X-Phi is facing the threat of redundancy, for evaluating a 'straw man'. Cappelen (2015), for example, voices this criticism of X-Phi with gusto:

> "[A] strong case can be made that Weinberg (and other experimental philosophers) are guilty of the most spectacularly uncharitable readings of philosophical texts in the history of Western philosophy. The entire movement is based on reading 2000 years of philosophy as relying on (or presupposing) a ridiculous (and obscure) view about something called 'intuitions'" (p.23).

So, setting the record straight on centrality is important.

1.2.2 What are Intuitions?

By 'intuition-talk', I mean instances of words like:

- Intuition
- Intuit
- Intuitive
- Unintuitive
- Counterintuitive
- Intuitively
- Counterintuitively

Few would deny that such terminology appears prolifically across contemporary analytic philosophy. Indeed, even most anti-centralists concede this, though they don't, of course, interpret it as citing evidential sources. Cappelen (2012), for example, admits that philosophers "undeniably" engage in quite a bit of intuition-talk (p.5), while Deutsch (2010) says:

> "How many times have we heard that a philosophical theory must be rejected because it has 'counterintuitive consequences'?" (p.448).

There's growing interest among philosophers in what this intuition-talk refers to, and that interest crosscuts the debate about the truth of centrality. On the one hand, if centrality is true, and 'intuition' refers to an evidential source, then just what is it that philosophers are citing in support of their claims? After all, to reiterate, settling this exegetical matter is a precondition to answering the more substantive questions, C and D, about whether our methods are any good; we need to know what this intuition-talk is referring to before we can assess philosopher's claims about what intuitions people have, and the usefulness of said phenomena as evidence for philosophical theories. As Smith (2000) says:

> "we cannot be sure whether we have such things, let alone what role they play in providing moral guidance[8], until we know precisely what intuitions are" (p.23–24).

On the other hand, if centrality is false, and intuition-talk isn't, in fact, being used to cite evidential sources, it's so common in our research papers that there would still be motivation to get clear about what work it *is* doing there. If this language is as important to our methods as its prevalence suggests, and its role has thus far been widely misinterpreted, there might be a

[8] While Smith focuses on the role of intuitions in *moral* philosophy, specifically, his claim is equally applicable to any area in which we find substantial intuition-talk (such as those areas of epistemology that I'll be looking at in this thesis).

different, important function that it's serving, yet to be discovered, or at least commonly acknowledged.

Furthermore, even if the emerging interest in intuitions and intuition-talk is primarily motivated by burgeoning metaphilosophical concerns, intuitions are of interest for reasons beyond that, rather parochial, academic curiosity. Intuitions are an interesting and important phenomenon in their own right, being a somewhat mysterious mental state that has hitherto resisted our attempts to conclusively analyse it, and something that apparently shapes and motivates many of our everyday beliefs. How often do we hear people apparently justifying their views by appealing to what we might generally call 'intuition', as in sayings like 'I just have an intuition' or 'I have a hunch'? Indeed, it seems we'll often even advise people to base beliefs on intuition, or something in its proximity, telling them things like 'trust your intuitions', 'trust your gut' or 'follow your instincts'. While it's plausible that philosophers use intuition-talk as a peculiar term-of-art, and therefore somewhat differently from its colloquial use among the general population, the languages likely have a shared etymology, and we should expect their meaning and use to overlap to some degree. So, saying something about philosophical intuitions specifically will still contribute to a larger picture of what intuitions are more generally, leading into bigger questions about what role they can and should play in the mental lives, not only of philosophers, but for all of us. Answering question B therefore seems a worthwhile pursuit in its own right.

1.3 Project Methodology and Novelty

In this section, I explain the methodological shortfall in the existing research into questions A and B, and how my own approach differs.

The single biggest shortfall in the existing literature regarding A and B is that, hitherto, metaphilosophical discussions about centrality have, for the most part, focussed on bold claims about the analytic practice in general, or otherwise been too broad in their approach. Concerning question A, PWI, for example, extrapolates generalisations about the role of intuitions in analytic philosophy as a whole – or lack thereof – from a case study of only a dozen seminal research papers. Likewise, the many rival answers to question B, as we'll see in the following chapters, present their views as theories about the nature of philosophical intuitions in general. There are, of course, some exceptions to this, as I'll point out as we go: in Ludwig and Sosa, the restriction of their account to intuitions cited in conceptual analysis, and, in Bengson and Brogaard, the identification of only one of (potentially) many senses of the word 'intuition' in the analytic tradition. Nonetheless, even those theories are meant to apply to overly broad swathes of philosophical research, such as to conceptual analysis in general, and the likes of Brogaard still

consider the sense they identify to be *the* dominant one across all analytic philosophy. These existing approaches are all, therefore, at odds with the heterogeneity of centrality, which, as I'll now explain, comes in two forms: diversity of application and ambiguity of language[9].

1.3.1 Diversity of Application

There are many different senses in which something can be used 'as evidence'. Consequently, there are many possible versions of centrality (the taking of intuitions 'as evidence')[10]. Indeed, as I'll now argue, there's reason to think that many different forms of centrality are being practiced in the analytic tradition.

One kind of centrality at work is a version of 'reflective equilibrium' (or RE). According to RE, philosophers should endeavour to achieve maximal coherence between their particular beliefs, both philosophical and not, as well as between those and the various principles, maxims, heuristics, *et cetera* that they use to decide them. Roughly, the idea is to continually adjust these elements until a sufficient 'equilibrium' is achieved. RE is most strongly associated with John Rawls (1971), who coined the phrase. Although Rawls did not, explicitly, do so himself, many have included 'intuitions' among the data to be reconciled. Broadly speaking, if a claim contradicts one of our intuitions, if it's 'counterintuitive', then one or both (though usually the former) should be adjusted to make them more coherent (see, for example, Goodman (1965, p.66)). Under this, 'intuition-inclusive' RE, intuitions are treated as evidence at least in the sense that they're taken to count against believing other things that contradict them. Call this 'RE centrality'.

---

[9] As mentioned here, we'll see more specific examples of anti-centralist arguments, and theories of intuition, making broad, unmitigated claims about the use of intuitions and 'intuitions' in philosophy *as a whole* as we go along.

[10] Intuitions might also play an important methodological role only insofar as they're the source of philosophical ideas, even if they're not taken to provide any evidence or justification for them. Bengson (2014, p.576) calls this the 'prompting' role of intuitions: an intuition that *p* 'prompts' one to reflect upon (the possibility that) *p*. Some, for example, read this practice into Kripke. Deutsch (2010), for instance, says ''the causal source of Kripke's judgment about the Gödel case is intuition; this much is fairly clear'' (p.451), and acknowledges that when Kripke calls *p* 'intuitive', he means to say that having a mental state (called an 'intuition') with content *p* was the causal origin of his idea that *p*. However, he denies that Kripke uses this fact as his evidence for *p*, in any sense; instead, Kripke proceeds to provide *arguments* for *p*. A similar practice is taking intuitions as 'working assumptions', or theoretical starting points. Here, an intuition isn't actually assumed, even defeasibly, to be true (by the author), nor is it used as evidence (for its own content or for something else). Rather, the intuited proposition is treated as a *hypothetical* truth. Then, one works through the entailments of the intuition, and concludes with either accepting or rejecting it, depending on how reasonable the outcome is (see, for example, Harman, 1968, p.166). This is similar to the sense in which one might 'assume' something is true in order to demonstrate a *reductio ad absurdum* by deducing a contradiction from it.

Another way of using intuitions as evidence is treating them as 'data points' to be explained by a theory. Here, the best theory is the one that can account for the most intuitions, *ceteris paribus*. Indeed, one formulation of centrality, from Pust (2000), frames the practice in exactly this way:

> "A theory is proposed... and it is tested by its ability to *account for* intuitive judgments" (p.5, my emphasis).

Call this 'explanation centrality'. To clarify, one can favour the theory that explains the most intuitions – the theory with the most 'explanatory power' – without also thinking that intuiting *p* counts against believing other things that are inconsistent with *p* (*a la* RE centrality), or that intuitions are a reliable guide to what's true, or a source of justification for (believing) their content. Indeed, one could simply explain intuitions by saying that they're, even systematically false and offering a plausible account of how those errors occur (an 'error theory'). Likewise, one might practice explanation centrality without ever using the propositional contents of intuitions as premises of an argument.

There are also some examples of intuitions being treated as 'evidence' in a more straightforward sense of the word, namely, where they're simply trusted to be a reliable guide to what's true. Goldman (1986), for example, claimed that the *correct* norms of belief are those that conform to the sense of justification that's "embraced by everyday thought and language" (p.58), and that we should judge which ones those are by their conformity to our intuition. Accordingly:

> "A criterion is supported to the extent that implied judgments accord with such intuitions and weakened to the extent that they do not" (*Ibid.*, p. 66).

Here, intuitions are treated, not just as something for our other views to cohere with (*a la* RE centrality), or as data points to be accounted for (*a la* explanation centrality), but as a reliable guide to what is, in fact, the case. Call this 'simple centrality'. A practitioner of simple centrality might treat intuitions as infallible (such that, if S intuits *p*, then *p*), or merely reliable to some degree (such that, if S intuits *p*, then *p* is at least likely, *ceteris paribus*).

We also have to be careful of a subtle distinction between 'intuitings' and 'intuiteds'. To elaborate, there are two, quite different ways in which 'intuition' can be read: either as the mental state itself (the intuit*ing*), or as the propositional content of that state (the intuit*ed*)[11]. This dichotomy has already been observed by, at least, Lycan (1998), Williamson (2000), Nado (2015), and Ichikawa (2014). The latter summarises it thusly:

[11] This obviously assumes that intuitions are mental states. In those cases, if any, where intuition-talk isn't even used to refer to psychological phenomena, the distinction mentioned here doesn't apply.

> "[W]hen Dave Chalmers has the intuition that zombies are possible, we can use 'intuition' to refer to a state of Chalmers's mind, as in 'the intuition explains his attraction to dualism'—or we can use it to refer to the content of this state, as in 'the intuition entails the falsity of physicalism'" (p.3).

Here, the 'state of Chalmers' mind' – his mental state, or the fact that he has it – certainly doesn't entail the falsity of physicalism, since his intuition could just be mistaken, but it might explain why that view appeals to him, while the propositional content of his intuition, 'philosophical zombies are possible', does entail the further proposition 'physicalism is false' (even if both claims are untrue).

We can therefore differentiate between at least two different metaphilosophical claims: the claim that philosophers rely evidentially on intuitings (mental states, or facts about them), and the claim that they rely evidentially on the intuited (the propositional contents of those states). Call the latter family of practices 'content centrality', because our evidence is the contents of intuitions, and the former 'psychological centrality', because we take as our evidence the mental states themselves, or facts about them, not their content. These methods aren't mutually exclusive, and can appear in conjunction with one or more of the other forms of centrality identified earlier, depending on in which sense the mental states and/or their content are employed as 'evidence'.

To help clarify the distinction here, we might make a comparison between content centrality and certain views of perceptual justification, which say that we take 'experiences' as evidence for *p* without inferring *p* from the experience itself, or some fact about it (that is, without psychologising our epistemic basis). Bengson (2014) has actually made such a comparison:

> "The basic idea is simple. Just as we need not think of ourselves as inferring beliefs about our immediate environment from psychologistic premises about our experiences in order to think of ourselves as relying on our experiences for positive epistemic status, we need not view Kripke or anyone else as making inferences from psychologistic premises about their intuitions in order to think of them as relying on intuitions for positive epistemic status" (p.570).

In other words, just as I might be relying on the content of my experiences when forming and justifying my beliefs about the world without basing anything on the experiential state itself, or the fact that I have it, the same might be said for intuition.

There's evidence of both content and psychological centrality being practiced in the analytic tradition. On the one hand, some seem to take content centrality as the obviously correct

interpretation of, at least, their own methods, especially when compared to psychological centrality. George Bealer (1998), for example, says:

> "[W]hen I say that intuitions are used as evidence, I *of course* mean that the *contents* of the intuitions count as evidence'' (p.205, emphasis added).

Here though, we must be cautious of yet another distinction. Specifically, we must consider the difference between cases where what is taken as evidence is merely coincidentally the content of an intuition, and cases where that content is taken as evidence *because* it's the content of an intuition. To elaborate, on the one hand, content centrality – the taking of intuited propositions as evidence – is compatible with it being entirely coincidental that an evidential proposition is intuited, it having nothing to do with *why* someone used that proposition, specifically, as evidence. I might intuit the factivity principle (if S knows *p*, then *p*), for example, and use that principle as evidence for the philosophical thesis that not all beliefs are knowledge (because some are false), but not because I intuited it; I would have used this principle regardless, so intuitions are methodologically incidental. Call this 'coincidental content centrality'. As Ichikawa (2014, p.234) has also observed, this form of the practice wouldn't, *ipso facto*, grant an interesting methodological role to intuitions. On the other hand, one might rely evidentially on intuited propositions *because* they're intuited. To clarify, it's possible for one to take the intuited content, *p*, as evidence because it's been intuited, still without actually citing the intuiting itself, or the fact that *p* was intuited, as evidence for *p*, or for something else. Again, we might compare this to perceptual experience: I might take the content of my visual experience, *p*, as evidence for something else, and consider myself warranted in doing so precisely because *p* was perceived, without ever citing the perception itself as evidence for anything. Here, that I experienced *p* isn't incidental to my taking *p* (not the experience that *p*) as evidence. Likewise, it might be that, while at no point is the intuiting itself cited, neither is the intuiting incidental to the selection of one's evidence. Call this 'warranted content centrality'.

Meanwhile, although psychological centrality has been called a misleading description of philosophical method by some, for allegedly 'over-psychologising' our evidence (Williamson, 2007; Deutsch, 2010), there are some areas of enquiry where such an approach seems uncontroversial, even obvious; specifically, those areas that deal with psychological questions about the mind and mental states themselves (see, for example, Goldman (2007)). For example, when philosophers are analysing intuitions, then they might indeed, and legitimately, take intuitions themselves as evidence for various theses. If, for instance, I intuit that *p*, when in fact *p* is false, then I might take the existence of that intuition itself, not its content, as evidence against the infallibility of intuition.

The above are just a few examples of how intuitions get used as evidence in the analytic tradition. We'll see some further instances of these methods in practice in the next chapter, as well as several other versions of centrality.

So, there are likely many different ways of using intuitions as evidence in philosophy. As I'll next explain, there's also, in all probability, a great deal of variation in what philosophers mean by said 'intuitions'.

1.3.2 Ambiguity of Language

When one looks at the range of answers proposed to question B, one thing that stands out about intuition-talk in philosophy is the total lack of consensus over its meaning. Most theories paint it as referring to some kind of mental state – an 'intuition' – but, even among those, there's tremendous variation in what kind of state is, allegedly, referred to. As Nado (2015) puts it:

> "It's difficult to isolate a single feature of intuition that elicits unanimous agreement—some philosophers deny that intuition has a special phenomenology, some deny that the justification it generates is a priori, and some even deny that it is immediate or unreflective" (p.785).

In addition, there's a growing number of theories that interpret intuition-talk as having some other, more pragmatic function, besides citing mental states, such as the view that it's simply a form of 'hedging' one's claims (Cappelen, 2012; Andow, 2017), or purely stylistic (Cappelen, 2012).

The level of disagreement here is virtually unprecedented. For any philosophically interesting phenomenon, there will inevitably exist much debate about its nature and properties, and we can imagine that for any claim made about it, no matter how commonsensical, there will always be someone, somewhere who denies it. Yet, there's usually at least some degree of consensus. For example, beliefs are an important and philosophically interesting mental state that have been subjected to centuries of analysis, and from that debate has emerged many competing theories of what beliefs are, what roles they play, and so on. However, it's almost unanimously assumed that beliefs are, at least, mental states, with content, that can be false. Even if some people deny one or more of these claims, even they must acknowledge that this is a shared core of assumptions in the literature. With intuitions, by contrast, there's no comparable level of consensus hitherto arrived at. It's still very contested whether intuition-talk even refers to mental states (as we've just seen), or whether such states are fallible (as we'll see later).

The peculiar degree of variation in how intuition-talk has been interpreted suggests one of two things may be going on. First, intuition could just be a massively nebulous concept in philosophy: such a hazy, unspecific term that it invites a plethora of conflicting interpretations. Even more cynically, we might say that, consequently, intuition-talk is semantically vacuous, failing to specify *any* particular referent at all. The latter has been argued most famously by Cappelen (2012), but a similar thought also appeared earlier, in Smith (2000), who likewise worried that our inability to provide a clear definition of 'intuition' might hide the fact that the term actually has no referent:

> "One suspects that the absence of definition, keeping intuition afloat as a hazy "something" between a thought and a feeling, may hide the fact that there are no such things" (p.23).

Alternatively, we could say that there's just a significant variation, rather than imprecision, in how intuition-talk is used in the literature, providing sufficient textual evidence to render plausible many different readings, which only go astray insofar as they overgeneralise. That would also explain why no consensus whatsoever has been achieved over even a single property of intuition: there's just no property that's common to everything that philosophers mean by 'intuition', because they mean many different things. Something resembling this view appears in Fischer and Collins (2015), who identify centrality as a 'theme'. Specifically, they identify the theme of contemporary analytic philosophy as:

> "Philosophers elicit, invoke, assess, and synthesize intuitions" (p.9).

This 'theme', they say, offers little descriptive content of actual philosophical practice, however, because it captures a range of very different and diverse, even contradictory, methods. All that unites those methods – or that creates the misleading impression of such unity – are common terms like 'intuition', but those terms are actually interpreted and employed in many different ways across the field.

I think we should favour this latter approach to intuition-talk, according to which the language is just ambiguous, in lieu of categorical evidence to the contrary, simply because it's more charitable[12].

[12] While I take it this is more forgiving than saying intuition-talk is semantically vacuous, it may still entail that such language is defective in academic contexts. If there's no standardised meaning for it to 'default' to, and no uniform properties of what we call 'intuition' that we can take for granted, then we have no way of telling what, specifically, a philosopher is referring to on a given occasion, and they may even be speaking of different things at different times, if they offer us no explicit definition (and they almost never do). Thus, even though I later identify what contextualists, specifically, mean by 'intuition', it may still be a defectively obscuring practice for them to call it that, in light of the ambiguity of such terminology.

As a direct consequence of this diversity of application and ambiguity of language, we have every reason not to generalise our metaphilosophical claims. There's simply nothing we can say about the meaning of 'intuition' in general, contra existing attempts to answer B, because philosophers likely use that term to refer to many different phenomena. Similarly, we can't pass general judgement on whether intuitions are used as evidence, and whether that practice is sound, since, to reiterate, there are many different senses in which intuitions might be 'used as evidence'.

This also has knock-on implications for addressing the aforementioned evaluative questions, C and D. For instance, we can't direct our epistemic critiques of intuition – our answers to D – toward that phenomenon broadly construed, because we'd have to dilute 'intuition' so much, to capture all of its senses, that we end up applying our conclusions to an absurdly broad range of phenomena. We actually find a whiff of this thought in Williamson (2007). Williamson likewise claims that there seems to be nothing consistent in how intuition-talk gets used: it merely stands for whatever it has to at the time to make the argument work. He criticises negative X-Phi on this very basis, saying that, because that research focuses on such a generic, nebulous term – 'intuition in general' – any subsequent scepticism about such 'intuitions' in philosophy that the research generates must be extended to most of our everyday judgements as well, given that the term refers to nothing specific – and hence nothing specific to philosophy – leading to what he calls 'judgement scepticism'. This amounts to, essentially, a *reductio ad absurdum* of existing experimental scepticism about intuition.

In the next section, I'll explain how my own research circumvents the issue of overgeneralisation by focussing on a very narrow, yet highly significant domain of philosophical literature.

1.3.3 My Focus: Epistemic Contextualism

To reiterate, in the formulation of centrality, words like 'intuitions' and 'evidence' should be read as umbrella terms for capturing a very wide range of phenomena and practices. The former can, arguably, stand for many different things, and those things can, demonstrably, be used in a variety of different ways. Due to combinatorial explosion, there are, therefore, a large number of potential philosophical methodologies that give an important role to 'intuitions'. Indeed, when he presented his formulation of centrality, Cappelen (2012) originally described it as a label for a family of theses, the specification of which depends on how we interpret its key terms (p.7).

Accordingly, my own, novel approach to centrality will be to focus on a very narrow domain of philosophical literature, in order to identify, through a careful, conscientious case study of that area, at least one thing that philosophers call 'intuitions', and how those things, in particular, get used in that domain. In other words, rather than coming up with some dubiously oversimplified

conclusions about philosophy as a whole, I'll look to some specific ways in which intuitions and intuition-talk get used, in the hope of saying something useful about, at least, one set of methodological practices at the core of one of the most influential and debated arguments in contemporary epistemology. This isn't to assume that the heterogeneity of centrality vanishes when we look more narrowly at specific areas (as we'll see in the next chapter, it certainly doesn't). It's just that, by narrowing the field of enquiry, we make the task of examining the various incarnations of centrality more manageable.

While the most immediate relevance of my findings will, therefore, be to the area I choose to investigate, my research will certainly be of more general significance to philosophy, foremost because the targeted literature, and hence the methodology I'm examining, has had considerable influence on the discipline as a whole. Thus, an analysis of this area reveals insights into how intuitions, of a particular variety, have been influencing our field more broadly. Also, by identifying and clarifying one of the phenomena we call 'intuition', and its role in our arguments, my work will contribute to the larger picture of our practices. Finally, by exercising a novel way to perform such metaphilosophical research, I'll also provide a model for this more considered approach to the questions surrounding centrality.

All that being said, the area of the literature that I've chosen to focus on will be the seminal arguments for epistemic contextualism. In the next sections, I outline this topic, and then my rationale for choosing it as the focal point of the investigation.

1.3.3.1 What is Epistemic Contextualism?

A word is semantically context-sensitive when its referent changes relative to the context in which it's uttered. Exactly how pervasive a phenomenon this is, in the English language specifically, is a subject of lengthy deliberation, but the simplest and least contentious examples are indexicals such as 'I' or 'here'. These words refer to different people or locations, depending on, at least, who utters them, and where, respectively. Consequently, changes in those specific contextual factors can, in turn, affect a change in their referents. The semantics of indexicals is, therefore, context-sensitive.

Broadly construed, epistemic contextualism (or EC) is the view that knowledge-attributing and knowledge-denying sentences – of the forms 'S knows that *p*' and 'S does not know that *p*', respectively – express different propositions (or propositions with different truth-conditions) depending on the context in which they're uttered, because the word 'knows', and its cognates, are context-sensitive, referring to different epistemic states relative to the context in which they're spoken. In other words, which epistemic state S must stand in for 'S knows that *p*' to come

out true – and, therefore, what S' epistemic position with respect to *p* must be to satisfy the truth-conditions of that assertion – vary relative to the context in which the knowledge is attributed. Prominent advocates of this position, or 'contextualists', include DeRose (2009, 2017), Cohen (2005), Lewis (1996), and Blome-Tillmann (2009, 2013). I'll explain their view in more detail in the next chapter.

1.3.3.2 Why Choose Epistemic Contextualism?

My research will identify specific iterations of centrality through a rigorous, comprehensive case study of the seminal arguments for contextualism. There are five reasons (i-v) I've chosen to focus on this area of the literature in particular.

i. *Prima facie*, the arguments for EC are a particularly fertile terrain in which to locate a centralist methodology, as they appear to rely, heavily, on appeals to intuition. Cohen (1986), for example, talks of 'the intuitions marshalled in support' of his own contextualism (p.583), and DeRose (1999), in his oft-cited paper *Contextualism: An Explanation and Defense*, uses intuition-talk no fewer than 22 times on 26 pages, which isn't even counting the multiple other phrasings that he also employs synonymously with words like 'intuition'. Indeed, there are two primary arguments for EC, and both seem to revolve around intuitions. The first is that the view can reconcile people's apparently conflicting intuitions about certain thought experiments, most notably DeRose's famous 'Bank cases' (1992), and Cohen's 'Airport cases' (1999). The second, which will be the focus of my own work, is that it can best explain the 'intuitive paradox' that arises when we consider intuitive knowledge claims and the, at least equally intuitive, closure principle for knowledge (now known, in some circles, as the 'Harman-Vogel paradox'; see, for instance, Nagel (2011)). Versions of this second line of argument appear in, for example, seminal works from DeRose (1995, 2017) and Cohen (1986, 1998, 2005). Consequently, this is a perfect place to investigate centrality at work in contemporary analytic philosophy.

ii. There's abundant evidence that the dominant interpretation of this literature is that contextualists do, indeed, practice centrality. For example, most critical discussions of EC frame the core challenge as being to offer an explanation and validation of people's 'intuitions', and present rivals to EC as potentially superior explanatory theses for those phenomena. See, for example, this representative abstract from Brown (2005):

> "Contextualism is motivated by cases in which the intuitive correctness of a range of phenomena, including knowledge attributions, assertions and reasoning, depends on the attributor's context. Contextualists offer a charitable understanding of these

> intuitions… Here, I investigate a range of different invariantist accounts and examine the extent to which they too can offer a charitable account of the contextualist data."

Meanwhile, a swathe of X-Phi endeavours to ratify or debunk what its practitioners are happy to call the 'intuitions' cited in defence of EC, and/or the empirical-psychological claims that contextualists make about what people intuit (Stitch, 2013; Sripada and Stanley, 2012; Phelan, 2013; Alexander *et al.*, 2014; Buckwalter and Schaffer, 2015), further testifying to the assumed reliance on intuitions, and facts about them, in that literature.

Because of how widely a centralist methodology is read into this literature, it's exactly the kind of place where discovering the falsity of centrality would have a serious impact on how we understand and engage with a seminal philosophical argument. Likewise, because there's so much X-Phi focussed on the 'intuitions' being appealed to in defence of EC, this is where the falsity of the centrality assumption would (potentially) have the biggest consequences for the experimentalist turn in philosophy. So, this is a very important place to get right about question A. Indeed, even if the answer to A is in the affirmative – which, as I'll argue in the next chapter, it is – identifying which, particular forms of centrality are actually at work here is crucial for any subsequent critique of that practice, as previously explained.

iii. EC is a particularly influential thesis in modern epistemology. According to a PhilPapers survey, it's become the dominant position (40.1%), relative to its antithesis, invariantism (31.1%)[13]. Indeed, the debate over its accuracy has since attracted many prominent figures in modern analytic philosophy, including Williamson (2005), Stanley (2004, 2007), and Hawthorne (2004, ch.4). Moreover, many modern alternatives to EC are, themselves, indirectly shaped by it. 'New invariantism', a term often used to denote positions such as Subject Sensitive Invariantism (or SSI) (Hawthorne, 2004) and Interest Relative Invariantism (or IRI) (Stanley, 2007), is a reaction to arguments for EC that takes the same 'intuitive data' seriously, and strives to produce contrary, non-linguistic explanations for it. The influence of EC, therefore, goes far beyond just those who adopt (some version of) it: it's been recognised as a large enough threat to also considerably reshape invariantist epistemology. All of this means that the philosophical methods practiced by contextualists have, for good or ill, been extremely important to shaping modern philosophy, and learning something about them will represent significant progress in our metaphilosophical research. Crucially, this is so even if it transpires that the method practiced in this literature isn't representative of the field as a whole, and doesn't frequently manifest elsewhere in the discipline.

[13] https://philpapers.org/surveys/results.pl [Accessed 29.08.20]

iv. As mentioned, the argument for EC that I'll be focussing on is its alleged ability to explain the Harman-Vogel paradox. As we'll see, this deals with a paradox of paradigmatic epistemic intuitions: intuitions about the closure, safety, and sensitivity principles, about scepticism, and so on[14]. Moreover, the Harman-Vogel paradox itself is one of the oldest and most tenacious intuitive paradoxes in history. The debate between contextualists and invariantists merely represents the most recent incarnation of an enduring philosophical challenge. So, the intuitions at the heart of this literature are quite central to philosophy more generally. Getting a clearer picture about what those are, and what role they play in the discussion, will therefore be hugely informative.

v. Despite the assumed role of intuitions in contextualist literature, and the resurgence of interest in centrality, there's hitherto been no careful investigation of the role of intuitions in the arguments for EC. Rather, the meaning of the abundant intuition-talk in that area, and how exactly such intuitions are employed evidentially, has been largely presupposed, and, as we'll see in the next chapter, the real picture of what's going on in these arguments is actually very complicated.

For these five reasons, I've chosen to focus my metaphilosophical investigation on the seminal arguments for EC.

1.4 Thesis Plan

I'll close this chapter by outlining my plan for the thesis, offering a brief synopsis of each subsequent chapter.

In Chapter 2, I explain EC in more detail, and the seminal argument for that thesis that I'm going to investigate, namely 'the argument from anti-scepticism'. Then, focussing on the works of, by far, the two most prolific and influential contextualist authors, namely Keith DeRose and Stewart Cohen, I carefully identify the multiple forms of centrality present in that argument, making the case for a centralist reading of this literature. I also provide counterarguments to anticipated anti-centralist critiques of such an interpretation.

In Chapter 3, I argue that extant minimalist accounts of philosophical intuitions fail to capture the intuitions being evidentially relied upon in the argument for EC, specifically. More exactly, I argue that said intuitions aren't outright beliefs, partial beliefs, dispositions to believe, or facts about ordinary language.

[14] I address potential concerns over whether the argument for EC really involves appealing to epistemic, as opposed to linguistic intuitions in Chapter 2.

In Chapter 4, I develop and defend a view according to which said intuitions are, instead, 'intellectual perceptual seemings'. However, I also argue that existing intellectual seemings theories of philosophical intuition are too conservative to capture the relevant phenomena, and propose a more liberal version of this thesis.

In Chapter 5, I argue that intellectual perceptual seemings, and thus the relevant intuitions, aren't a *sui generis* phenomenon, as many are wont to assume. Rather, they're a sub-category of metacognitive experiences that some psychologists call 'epistemic feelings'.

Finally, in Chapter 6, I summarise my findings, and briefly explore some of their implications. In particular, I consider the possibility that we should treat certain intuitions as higher-, not first-order evidence for their content.

Overall, then, I conclude that certain intuitions, and thus certain 'epistemic feelings', have played a significant part in shaping modern epistemology, and that this revelation has significant ramifications for how metaphilosophical investigations of that domain should proceed.

Chapter 2 Epistemic Contextualism and Intuitions

In this chapter, I argue that something called 'intuitions' play a central, evidential role in one of the primary arguments for epistemic contextualism, namely 'the argument from anti-scepticism'. In later chapters, I'll defend a novel account of what said intuitions are.

In §2.1, I outline the contextualist thesis. Then, in §2.2, I explain 'the argument from anti-scepticism'. In §2.3, I present my case for a centralist reading of this argument, identifying several ways in which it treats intuitions as 'evidence', broadly construed. Finally, in §2.4, I respond to both existing and anticipated anti-centralist objections to such an interpretation of philosophical method.

2.1 Epistemic Contextualism

In this section, I explain the view known as epistemic contextualism, expanding upon the brief outline of that position that I offered earlier.

As previously explained, a word is semantically context-sensitive when its referent changes relative to the context in which it's uttered. The simplest and least contentious examples of this phenomenon are indexicals such as 'I' or 'here'. These words refer to different people or locations, depending on, at least, who utters them, and where, respectively.

Broadly construed, epistemic contextualism (or EC) is the view that knowledge-attributing and knowledge-denying sentences (hereafter 'K-sentences', following DeRose, 2000) – of the forms 'S knows that *p*' and 'S does not know that *p*', respectively – express different propositions (or propositions with different truth-conditions) depending on the context in which they're uttered, because the word 'knows', and its cognates (hereafter 'K-words'), are likewise context-sensitive, referring to different epistemic states relative to the context in which they're spoken. In other words, which epistemic state S must stand in for 'S knows that *p*' to come out true – and, therefore, what S' epistemic position with respect to *p* must be to satisfy the truth-conditions of that assertion – vary relative to the context in which the knowledge is attributed.

Advocates of EC, known as 'contextualists', include, most notably, DeRose (2009, 2017), Cohen (2005), Lewis (1996), and Blome-Tillmann (2013). The antithesis of this view is 'invariantism', and thus its critics are known as 'invariantists', from terminology coined by Unger (1975). Invariantists maintain that, contrary to EC, the referents of K-words don't vary with context.

Unlike indexicals like 'I' or 'here', the referents of K-words aren't meant to be fixed by basic environmental factors, such as who is uttering them or where, but rather by the *epistemic standards* that obtain in their context of use. To clarify, contextualists sometimes compare K-words, not to indexicals like 'I' or 'here', but to context-sensitive adjectives, like 'tall', 'flat', and 'cold'[15]. If Carla is 6 feet tall, then 'Carla is tall' expresses a true claim in most contexts. However, if someone is talking specifically about basketball players, then the same claim, 'Carla is tall', might instead express a falsehood. That's because 6 foot is tall *for most people* but not *for basketball players*. When someone is discussing the height of basketball players, they're likely using different standards for 'tallness' than when they're discussing people in general. What 'tall' means, in any given conversational context, is at least partly dependent on those standards. Likewise, '4 degrees is cold' might express a true claim when uttered in England, but not when uttered in Greenland, again because there are different standards of coldness in play, and hence 'cold' means something different in the two contexts. EC says something similar about K-words, which pick out different epistemic positions depending on the standards that apply to the current conversational context. However, while the standard for something like tallness is clearly some minimum height, and what determines that standard is something like the average height of the people-type one is talking about, rival versions of EC disagree over what *epistemic* standards actually are, which aspects of those standards are context-sensitive, and which contextual factors, specifically, they're sensitive to, as we'll see later.

EC's most contentious implication is, therefore, that a K-sentence, like 'S knows that *p*', might be used to express a true claim in one context, and a false claim in another, even if both utterances are made simultaneously, about the very same S and *p*. This will be the case whenever S satisfies the epistemic standards that obtain in the former context, but not those which obtain in the latter context. Likewise, the second speaker could instead utter 'S doesn't know that *p*', and both speakers would express true propositions in their respective contexts. The sentences would be 'surface contradictory' but, because 'knows', in one context, and 'know', in the other, are referring to different epistemic positions, they're actually compatible. As Cohen (1988) puts it:

> "[O]ne speaker may attribute knowledge to a subject while another speaker denies knowledge to that same subject, without contradiction" (p. 97; c.f. DeRose, 1992, p. 920).

Finally, we should clarify that, for EC, what matters is the context in which knowledge attributions are uttered, not the context in which the putative subject of that knowledge finds themselves. It's

[15] For an example of this comparison, see Cohen (1986, p.580).

almost platitudinous to observe that facts about the subject's context can affect whether they know something, and hence the truth-*value* of propositions like 'S knows that *p*'. Merely by changing facts about S' environment, we can change what's true, for example, and thus what it's possible for them to know. If, for instance, it's true that 'S knows they're looking at a barn', we can make that false by substituting the real barn for a fake one. However, 'subject factors' (facts about S and their context) need not affect the truth-*conditions* of K-sentences. So far, we could hold onto invariantism, maintain that the knowledge position picked out by K-words is always the same, and simply claim that S satisfies it before we switch the real barn for a fake one, and not afterwards. Facts about S's context affect the truth-*values* of K-sentences (in this example, whether factivity is satisfied), but not the truth-*conditions* of such claims (whether factivity is necessary). By contrast, EC is a thesis about the role that 'attributor factors' (facts about the speaker of the K-sentence, and their context) are playing. According to this, not only does the subject's context affect whether S satisfies a specific epistemic position, and hence the truth-values of K-sentences, but the attributor's context also determines precisely which epistemic position they are actually attributing to S, and hence the truth-*conditions* of their K-sentences[16].

In the next section, I'll explain one of the core arguments for EC, namely the 'argument from anti-scepticism'. Then, in the remainder of the chapter, I'll make the case that said argument embeds multiple forms of centrality, and conclude that intuitions have, therefore, played a crucial role in the rise of contextualism.

2.2 Contextualism and Scepticism

In this section, I'll explain one of the primary arguments for epistemic contextualism (or EC), namely the 'argument from anti-scepticism'. This argument essentially claims that we can only provide a satisfying reply to classic, closure-based sceptical arguments by adopting some form of contextualism. As DeRose puts it:

> "Contextualist theories of knowledge attributions have almost invariably been developed with an eye toward providing some kind of answer to philosophical skepticism" (DeRose, 2017, p.3).

[16] Of course, subject factors and attributor factors will overlap when speaker and subject are the same person, such as in cases of first-person knowledge attribution, or when they otherwise share the same context. Hence, subject factors may indeed affect the referents of K-words, but only in virtue of also being attributor factors.

Indeed, DeRose (2017) devotes almost the entirety of his recent book – *The Appearance of Ignorance: Knowledge, Skepticism, and Context, Vol. 2* – to developing this argument, providing us with the most comprehensive, state-of-the-art version of this case for EC. Our other eminent contextualist author, Cohen, has also pursued this line of argument extensively in several of his seminal papers (1988, 1998, 2005).

I'll begin by explaining the sceptical problem itself, followed by the proposed contextualist solution.

2.2.1 Closure-Based Scepticism

As the argument for EC rests on its ability to respond convincingly to closure-based sceptical arguments, we should first get clear about what those arguments actually look like, and the problem they create. The primary target of DeRose's recent book is what he calls 'The Argument from Ignorance' (or AI):

1a [Premise]: I don't know not-*H*

2a [Premise]: If I don't know not-*H*, then I don't know *O*

3a [Conclusion]: I don't know *O*

H denotes a sceptical *hypothesis*, one that we normally take ourselves to not know the negation of, and *O* a proposition that we *ordinarily* take ourselves to know, but which is incompatible with the truth of *H*.

Cohen (1988, p.93-94), meanwhile, discusses a very similar argument, though his resembles more closely the form of *modus tollens* than *modus ponens*:

1b [Premise]: If S knows *q* and S knows that *q* entails ~*H*, then S knows ~*H*

2b [Premise]: S does not know ~*H*

3b [Conclusion]: ~(S knows *q*)

S denotes an arbitrary epistemic subject, *H*, again, a sceptical hypothesis, and *q* a proposition that we ordinarily take ourselves to know, but which is incompatible with the truth of *H*. Note that Cohen's version of the sceptical argument, unlike DeRose's, isn't even valid; the negation of the consequent in premise 1b only entails that the antecedent conjunction is false, which does not entail that one particular conjunct, 'S knows *q*', is false, rather than 'S knows that *q* entails ~*H*'.

However, as I'll soon explain, it's the intuitiveness of sceptical arguments that contextualists want to address, not their soundness.

Both DeRose's and Cohen's arguments are designed to reflect the common structure of *all* closure-based sceptical arguments: they involve (in either order) an intuitive claim to not know the negation of some radical sceptical hypothesis, and an at least equally intuitive claim to the effect that, if we don't know that negation, then we don't know some other, intuitively known proposition either, entailing the counterintuitive conclusion that we don't, after all, know the latter.

The sceptic's argument doesn't, of course, claim that anybody really is the victim of a sceptical scenario. Instead, it argues that, even if S isn't in such a scenario, their inability to know that they're not (to know ~*H*) means they don't know *O* or *q*. Thus, AI-style arguments seem to threaten our ordinary claims to knowledge (our claims to know *O* or *q* propositions). They can also threaten us with a broader, more all-encompassing scepticism, for two reasons. First, because, as stated, *O* or *q* will be a paradigmatic example of what we ordinarily take ourselves to know, and hence something that we consider ourselves to be particularly well-situated to know, such that, if we don't even know *that*, then it seems we don't know much else either, if anything. And second, because, for any *O* or *q*, one can identify a corresponding *H*, such that the argument can be repeated, *ad nauseam*, for every putatively known proposition, until we arrive at a more general scepticism.

It could be said, however, that the real worry triggered by sceptical arguments isn't that we might actually lack knowledge (in general, or of *O* or *q* in particular). Even if we assume that such knowledge is real, sceptical arguments still present us with a puzzle, specifically an 'intuitive paradox', known, in some circles, as the 'Harman-Vogel paradox' (Nagel, 2011). The general form of the Harman-Vogel paradox is as follows: S intuitively doesn't know ~*H* and, intuitively, if S doesn't know ~*H*, then they don't know *O* (or *q*) either but, intuitively, S *does* know *O* (or *q*). Or: S intuitively knows one thing, and not another, but it's at least equally intuitive that if S knew the former, then they'd know the latter as well. To put it simply, we have an inconsistent set of intuitions. The job of the modern epistemologist isn't so much to defeat the sceptic (after all, there are almost no advocates of scepticism in contemporary philosophy to argue against[17]), but rather to incorporate, into their broader theory of knowledge, an explanation for how such knowledge is possible, despite the intuitiveness of the sceptic's premises, and what burdens us

[17] Unger (1975) is a noteworthy exception.

with these paradoxical intuitions in the first place. As we'll shortly see, this is certainly how contextualists, at least, treat such arguments.

Henceforth, to help make matters less abstract, I'll discuss a specific example of a sceptical argument, rather than their general form. Let's take DeRose's formulation of such arguments as our starting point, simply because it's neater (and valid). However, let's substitute 'I' for an arbitrary epistemic subject, S, as Cohen does. For *H*, we can, as DeRose frequently does, use the claim that S is the victim of Putnam's (1981, p.5-6) classic brain-in-a-vat sceptical scenario, the details of which are probably quite familiar: S is a disembodied brain, floating in a vat of life-preserving fluids, attached to a supercomputer that stimulates them in exactly the ways they would be stimulated if they were walking around leading a normal life, thereby feeding them a subjectively indistinguishable experience as-if doing so. Let's therefore take, as our *H*, the proposition that S is a brain-in-a-vat (or BIV), and call the premise that S doesn't know the negation of this proposition 'ignorance'. As our *O*, let's use the proposition that S has hands (and stipulate that S does, indeed, have hands), and call the conclusion that S doesn't know this 'scepticism', owing to its aforementioned implications. Finally, let's call the bridging premise 'closure', for reasons that I'll explain shortly. This gets us:

Ignorance:	S doesn't know that they are not a BIV.
Closure:	If S doesn't know that they are not a BIV, then S doesn't know that they have hands.
Scepticism:	S doesn't know that they have hands.

In this instance of the Harman-Vogel paradox, S intuitively doesn't know they're not a BIV, but the, at least equally intuitive 'closure' carries us from that to the *counter*intuitive conclusion that S doesn't know they have hands either.

Henceforth, I'll refer to the above, in DeRose style, as the 'argument from ignorance' (or AI), because it likewise begins with a claim to ignorance.

2.2.2 The Alleged Data

As stated in section 2.2.1, closure-based sceptical arguments, like AI, present epistemologists with an 'intuitive paradox' (the Harman-Vogel paradox). As we'll see, contextualists defend EC on the grounds that it can explain away that paradox while avoiding scepticism. To really understand their explanation, however, it's necessary to first get clear about what they think needs

explaining, that is, what they take our intuitions – about 'ignorance', 'closure', and 'scepticism' – to be, in more detail.

First, 'ignorance'. DeRose considers this premise intuitive. For one, every piece of experiential evidence that S has for leading a normal life seems to be 'undermined' by the sceptical scenario in question, as does any empirical evidence they might possess pertaining to its scientific implausibility. As DeRose (2017) puts it, we have:

> "strong *intuitions* to the effect that the things we learn about the world through experience get 'undermined' if we try to use them as evidence for [~BIV]" (p.249, my emphasis).

There are several possible ways to flesh out this thought. We might think that, for instance, appealing to experiential evidence is question-begging against the sceptical hypothesis, since, if S really is a BIV, all of those experiences are misleading. Likewise, using the fact that the BIV scenario is implausible given S' current scientific understanding begs the question against the BIV hypothesis by assuming the veridicality of S' empirical observations as-of the universe being that way. Alternatively, we might say that S would still have their exact perceptual experiences as-of leading a normal life if they were, indeed, a BIV, making their belief, grounded in that evidential *corpus*, that they're experiencing the real world, 'insensitive' to the facts, and such insensitivity is incompatible with knowledge. Of course, one doesn't have to agree with either of these more specific claims, or even with the first premise of AI in general, to agree that 'ignorance' is, at least, intuitive.

Cohen, meanwhile, discusses sceptical hypotheses in general, and not our BIV intuitions specifically. He does, however, think that the claim to ignorance (his 2b) is plausible, where *H* denotes any familiar sceptical hypothesis (Cohen, 1988, p.94). We can therefore take Cohen to think that his 2b will seem plausible when *H* is substituted for being a BIV.

So, both our contextualists would agree that 'ignorance' is intuitive. However, DeRose thinks we can also find it absurd to suggest that, for all we know, such an imaginative, elaborate, far-fetched, science-fiction scenario as being a BIV might actually obtain. As he puts it:

> "To the suggestion of BIV, we are moved to exclaim: 'Absurd!', 'Outlandish!', 'Aw, come on!', 'Get out of here!', 'Come off it!', 'That's ridiculous!' To ~BIV: 'Of course!'" (2017, p.228).

So, for DeRose, while, at times, it might seem that 'ignorance' is obviously true, at others it might seem quite ridiculous. He also stresses how we can experience both of these 'intuitive pulls' simultaneously, leading to cognitive dissonance:

> "Though I could always feel a strong intuitive pull toward accepting this premise, I found it far from compelling, and I also felt a significant opposing intuitive push toward denying it" (*ibid.*, p.55).

Cohen, on the other hand, doesn't seem to (explicitly) endorse this further claim.

Next, 'closure'. Again, both our contextualists claim this premise is intuitive. For one, it seems to be supported by the, highly intuitive, 'closure principle for knowledge' (or CPK), which states that knowledge is closed under known entailment. In Cohen's version of the sceptical argument, his 1b – 'If S knows *q* and S knows that *q* entails ~*H*, then S knows ~*H*' – is simply his personal formulation of this principle, which he says strikes us as "compelling" (Cohen, 1988, p.93). DeRose (2017), on the other hand, uses a slightly different iteration, which can be translated as:

> If S knows *p*, and *p* entails *q*, then S knows *q*

This is very close to Cohen's version of CPK, only it doesn't require S to actually know the entailment. From this, we get the contrapositive:

> If S doesn't know *q*, and *p* entails *q*, then S doesn't know *p*

The latter is reflected in DeRose's 2a: if S doesn't know ~*H*, then S doesn't know *O*, because *O* entails ~*H*. DeRose likewise says that his iteration of CPK is intuitive, and thus says:

> "AI's other (second) premise is underwritten by 'closure' intuitions" (*ibid.,* p.54, Footnote 12).

Admittedly, DeRose and Cohen's formulations of CPK are both dubious principles. DeRose's entails that S knows all the implications of what they know, which is implausible. For example, *p* entails any, indefinitely long disjunction of arbitrary propositions (*p* or *q* or *r* or…*n*), and it's too strong to say that S must also know each of these disjunctions to know that *p*. Cohen's version is somewhat more restrictive, by requiring S to first know the entailment holds before they must know what is entailed by their knowledge, but this is likewise too strong a claim. S might know *p*, and that *p* entails *q*, but not even believe *q*; they might, for instance, never actually make the inference from these two facts to *q*.

However, as previously mentioned, it's the intuitiveness of the sceptic's claims that contextualists are trying to explain, not their accuracy. It can, for example, sound intuitive to say, *a la* DeRose:

> If I know I have hands, then I know I'm not a brain-in-a-vat.

Or, *a la* Cohen:

> If I know I have hands, and I know that having hands entails I'm not a brain-in-a-vat, then I know I'm not a brain-in-a-vat.

While false upon reflection, there's something initially compelling about these statements – perhaps because we would *expect* someone who utters them to know that having hands entails not being a BIV, and to have already made that inference. Similarly, while Cohen's formulation of CPK doesn't, in combination with 'ignorance', entail 'scepticism', as I previously mentioned, we might nonetheless find it intuitive that it does because we would expect someone to *at least* know that *q* entails ~*H* – to know that having hands entails that one is not a handless BIV, for instance – or to come to know that as soon as they start to consider those propositions, such that, we would *expect* someone who doesn't know the antecedent conjunction in Cohen's 1b to lack knowledge of the other conjunct ('S knows *q*') instead.

So, 'closure' is at least supported by intuitive, if not sound epistemic principles.

Closure-based sceptical arguments are so-called precisely because they appeal to CPK intuitions in this manner, and I've elected to call the connecting premise in AI 'closure' because of its intimate relation to that principle. However, that premise is also meant to be independently intuitive, no matter how we picture its relationship to CPK, or believe the correct formulation of that principle to be. DeRose (2017), for example, claims that that it will seem to most people that S is, at least, no better positioned to know that they have hands than they are to know that they are not a handless brain-in-a-vat:

> "[W]e can sense that… I am in no better a position to know that O than I am in to know that not-H… Closely tied to that… is the related and intuitively compelling realization that it would be no wiser to bet one's immortal soul on O's being true than to bet it on not-H's being true" (p.25).

So, one doesn't have to agree with either author's formulation of CPK to recognise that they, or something it their vicinity, is intuitive, or agree with CPK at all to appreciate the independent plausibility of 'closure'. For our purposes, what's important is just that both contextualists take the sceptic's second premise to be highly intuitive.

Finally, 'scepticism'. DeRose and Cohen both claim that, on the one hand, because *O* or *q* will be a paradigmatic example of what we ordinarily take people to know, the sceptic's conclusion (that we don't know *O* or *q*) will seem false. This certainly applies when *O* or *q* is, specifically, 'S has

hands'. In fact, DeRose quotes Moore (1959) on the absurdity of denying knowledge of this very proposition:

> "'Here is one hand and here is another': How absurd it would be to suggest that I did not know it, but only believed it, and that perhaps it was not the case!" (p.146).

On the other hand, however, because the premises of the sceptic's argument seem true, that argument seems sound, and we feel at least somewhat tempted to accept its conclusion. Cohen (1998), for example, says:

> "[W]hen in the grips of a sceptical argument, we feel the strong intuitive pull of saying that our knowledge ascriptions are mistaken" (p.302).

While DeRose (2017) speaks of "the intuitive power of the skeptic's claims" (p.59). So, when confronted with the sceptical argument, our intuitions about its conclusion become conflicted, and/or fluctuate (see, for example, Cohen (1986, p.582)).

In summation, our ambivalence toward 'ignorance', our confidence toward 'closure', and our confusion over 'scepticism' are the primary data points constituting the Harman-Vogel paradox, according to the contextualists. Next, I'll explain how EC proposes to resolve said paradox while avoiding scepticism.

2.2.3 The Contextualist Solution

Contextualists defend EC on the grounds that it can resolve the aforementioned Harman-Vogel paradox, while avoiding scepticism. In this section, I explain said resolution.

Contextualist treatments of sceptical arguments, like AI, although they vary in their details, are all fundamentally the same. The shared idea is that claims to know things like 'S has hands' come out true when uttered in most, everyday contexts, but that something about having the sceptical argument presented to you causes a shift in epistemic standards, such that the truth-conditions for such claims are now different, and more demanding; S fails to satisfy the epistemic position toward having hands now picked out by K-words, so the claim to know they have hands, when uttered in the new context, comes out false. Yet, this concession doesn't change the fact that everyday knowledge claims are true. Cohen, for example, summarises the contextualist approach, broadly construed, thusly:

> "The basic idea is this: the sceptic's appeal to hypotheses (involving brains-in-a-vat, evil demons, etc.) creates a context where the standards for knowledge… are stricter than the standards that govern typical everyday contexts. In those 'sceptical' contexts, we fail

> to know anything. So the contextualist concedes that there is some truth to scepticism. But a contextualist can limit the damage in a crucial way. For it remains true that we know many things in the typical everyday contexts where the standards are lower" (Cohen, 1998, p.291-292).

The particulars of the contextualist solution to scepticism vary according to how different contextualists understand the specifics of their view. Most crucially, contextualists disagree among themselves about three factors: epistemic strength, epistemic standards, and what causes said standards to change. In order to more fully understand how EC replies to the sceptic, and in particular how it accounts for our paradoxical intuitions about AI, we'll need to first explore some of these finer, more idiosyncratic details.

Epistemic strength is how good S' epistemic relation to a proposition, *p*, is. DeRose (2017), for example, understands this is terms of how 'safe', or truth-tracking, S' belief that *p* is:

> "The further away one can get from the actual world, while still having it be the case that one's belief matches the fact at worlds that far away and closer, the stronger a position one is in with respect to P" (p.26)[18].

Cohen (1988), meanwhile, cashes out epistemic strength in terms of how many alternatives to *p* (propositions that are mutually exclusive with *p*) are incompatible with (or 'eliminated' by) S' evidential *corpus*, where S' epistemic position with respect to *p* is stronger the fewer alternatives to *p* remain 'uneliminated' (p.94)[19].

Epistemic standards, meanwhile, are the minimum strength of epistemic position toward *p* required for knowledge of *p* in that context, and therefore determine (at least partly), the truth-conditions of knowledge ascriptions made in said context. Contextualists' views on this correspond to how they understand epistemic strength. DeRose (2017), for example, understands epistemic standards in terms of how far from the actual world S' belief that *p* must track the truth (or how 'safe' S' belief that *p* must be) for S to know *p*. For him, we can therefore picture epistemic standards as:

---

[18] There are two noteworthy details about this. First, DeRose employs a double-safety notion, not a single-safety notion, that is, S must not only avoid believing *p* in worlds where *p* is false, but also disbelieving *p* in worlds where *p* is true. Second, DeRose understands the 'closeness' of possible worlds specifically in terms of their (relevant) similarity.

[19] Later, Cohen (1998) clarifies by borrowing from Lewis' framework. Lewis defines an uneliminated alternative as one where the subject's evidence, understood as their "perceptual experience and memory", in that alternative scenario, "exactly match his perceptual experience and memory in actuality" (Lewis, 1996, p.553).

> "[A] contextually determined sphere of possible worlds, centered on the actual world, within which a subject's belief as to whether P is true must match the fact of the matter in order for the subject to count as knowing… Call this sphere the sphere of epistemically relevant worlds…" (p.28).

Conversely, Cohen understands epistemic standards in terms of the set of alternatives to *p* that S' evidential *corpus* must eliminate. That is, for him, epistemic standards are represented by a set of epistemically relevant propositions, not worlds (Cohen, 1988, p.94)[20].

Finally, in terms of what causes epistemic standards – and hence the truth-conditions of K-sentences – to change with the attributor's context, there's a great deal of variation among contextualists. Some have suggested that the matter is settled by how important it is for the speaker of the K-sentence to be right about *p*: generally, the more severe the consequences of being wrong about *p*, the stronger ones epistemic position with respect to *p* must be to know *p*. DeRose has toyed with this idea in a number of places (1999), as has Hambourger (1987, p.260), and Cohen (1988, p.97), who ties it, specifically, to justification: the more severe the consequences if ~*p*, the more evidence for or reason to believe *p* S needs for such a belief to be justified, and justified belief is a precondition of knowledge.

Others link epistemic standards to conversational manoeuvres. Lewis (1979, p.352; 1996, p.551-60), for instance, argues that merely mentioning new alternatives to *p* is sufficient to make them epistemically relevant to knowing *p* in that context, providing the raised possibility isn't ignored by the conversational participants. This resembles an earlier thought in Austin (1946, p.113), and is also similar to Cohen's (1986) more considered position, which is that mentioning new alternatives can change epistemic standards if it makes those possibilities 'subjectively evident', or salient, to the speaker of the K-sentence. To elaborate, for Cohen, if an alternative to *p* is sufficiently likely, objectively or conditional on the speaker's evidence, then it's epistemically relevant. Conversely, if an alternative is sufficiently improbable, objectively and conditional on the speaker's evidence, then it's not relevant. However, even if an alternative is extremely unlikely, drawing the speaker's attention to it (making it 'salient' in the conversation) means that this statistical fact is no longer sufficient to make it epistemically irrelevant. Indeed, it then becomes relevant. Blome-Tillman, meanwhile, thinks that mentioning possibilities only affects epistemic standards if doing so alters the speaker's 'pragmatic presuppositions' (2009, p.256). For example,

[20] Under some ways of understanding 'worlds' and 'propositions', some of these differences will be rendered purely cosmetic. On the view that propositions are just sets of worlds, for example, Cohen's relevant alternatives are just sets of relevant worlds. But some differences seem more fundamental. The view that epistemic standards require a belief to be truth-tracking across those worlds, for instance, seems quite different from the view that they require the *elimination* of those worlds.

normally, people implicitly presuppose that they're not a BIV, and so that possibility isn't epistemically relevant for them. But, if one raises the possibility of being a BIV, it might cause speakers to not assume that they aren't, and thus make that possibility relevant.

Finally, DeRose imagines several contextual parameters to which epistemic standards might be relative, but the one that is meant to explain the data about AI is which explicit knowledge claims have been made. Here, DeRose's approach makes considerable use of the principle of sensitivity:

> Sensitivity: S would not have believed that *p* if ~*p*

This just means that, in the closest world(s) to the actual world where ~*p*, S doesn't believe *p*. Crucially, DeRose (2017) says: "I do not take sensitivity to be a [general] requirement for knowledge" (p.116). That is, he doesn't take sensitivity to be a truth-condition for all assertions of the form 'S knows that *p*'. However, he does think that, in *some* contexts, a belief that *p* must indeed be sensitive to count as knowledge. To understand how, we have to appreciate how sensitivity is related to safety: if one's truth-tracking extends to the closest ~*p* worlds, then one's belief that *p* is sensitive. DeRose claims that epistemic standards (recall, how safe S' belief must be) can be raised to the point of sensitivity – that is, the sphere of epistemically relevant worlds can be expanded, if necessary, to encompass at least the closest ~*p* world(s) – by explicit utterances of claims to know, or not know, *p*. This is DeRose's crucial 'rule of sensitivity':

> "When it is asserted that some subject S knows (or does not know) some proposition P, the standards for knowledge… tend to be raised, if need be to such a level as to require S's belief in that particular P to be sensitive for it to count as knowledge…" (*Ibid.*, p.27)[21].

With these specifics and variations in mind, we can now grasp the contextualist response to scepticism, and its explanation of the aforementioned intuitive data, in more detail. Again, let's focus on, specifically, the versions of EC advocated by Cohen and DeRose.

First, both Cohen and DeRose can explain why 'ignorance' seems true when we hear it. For DeRose, because it's an explicit claim that S doesn't know they're not a BIV, it expands the sphere of epistemically relevant worlds to include at least the closest worlds where S is a BIV. As DeRose (2017) observes: "if I *were* a BIV, I would believe every bit as firmly as I actually do that I *wasn't* one" (p.14, emphasis in original). Consequently, S' belief that they're not a BIV doesn't satisfy the

[21] While the Rule of Sensitivity plays a pivotal role in DeRose's explanation of our intuitions about AI, he notes that: "This mechanism for raising the standards for knowledge is just one among many conversational forces affecting epistemic standards" (*Ibid.*, p.189).

new epistemic standards: it isn't safe enough. Thus, when presented with AI, 'ignorance' tends to become true (in your context). In other words:

> "The reason you can seem not to know that you're not a BIV… is that talk of whether you 'know' such things tends to invoke unusually high epistemic standards that you do not meet" (*Ibid.*, p.162).

Indeed, owing to the rule of sensitivity, "whenever S's belief that P is insensitive, we can *truthfully* assert that S *doesn't* know that P, and can only *falsely* say that S *does* know that P" (*Ibid.*, p.29, emphasis in original).

Meanwhile, for Cohen, we feel that 'ignorance' is true when we hear it because of his 'rule of salience': hearing or thinking about S being a BIV makes that a salient possibility, which makes it an epistemically relevant alternative to S *not* being a BIV (in your context). This adds it to the list of propositions that S' evidence must outright eliminate before they know they're not a BIV, despite the objective unlikelihood of them actually being one:

> "[W]hen we are confronted with skeptical arguments, we may come to consider skeptical alternatives as relevant, thereby lowering our standards (for how probable an alternative must be)" (Cohen, 1988, p.96).

Again, S cannot satisfy the new epistemic standards for knowing they're not a BIV: their evidence isn't mutually exclusive with that possibility. Thus, again, when presented with AI, 'ignorance' tends to become true (in your context).

Second, contextualists explain the intuitiveness of 'closure' by simply preserving its truth. Indeed, the conditional, that if 'ignorance' is true then 'scepticism' is true, holds regardless of the epistemic standards in play. Specifically, this claim is preserved by DeRose because, if there are no epistemically relevant worlds in which S falsely believes they have hands, then there are no such worlds where S falsely believes they're not a BIV either, for, if S falsely believed they weren't a BIV, they would also falsely believe that they had hands. Put simply: if S' beliefs track the truth about having hands well enough to know they do, then they also track the truth about not being a BIV well enough to know they're not. Thus, by *modus tollens*, if S doesn't know they're not a BIV, they don't know they have hands either. For Cohen, meanwhile, 'closure' follows because, whenever S' evidence is sufficient to know that they have hands, it must be sufficient to either eliminate the alternative that they're a BIV, if it's relevant, or to make such an alternative too subjectively improbable to be relevant. Cohen maintains that, if S' evidential corpus is sufficient to make S being a BIV epistemically irrelevant, then S can know they're not a BIV on the basis of that evidence:

> "*[I]n context c*, h is not a relevant alternative to not-h. That is to say… in c, e is sufficient evidence to deny h, to prevent knowledge of not-h from being precluded… Thus, e is sufficient evidence to believe not-h, to prevent knowledge of not-h from being precluded. It follows that e is sufficient evidence to know not-h" (Cohen, 1988, p.105, emphasis in original).

So, in any case where S knows they have hands, their evidential body is sufficient to either eliminate the possibility that they're a BIV, or to at least make it epistemically irrelevant, and either way S gets to know that this possibility doesn't obtain. Thus, again, by *modus tollens*, if S doesn't know they're not a BIV, they don't know they have hands either.

Third, both Cohen and DeRose can explain our confusion over 'scepticism'. On the one hand, S does know they have hands by ordinary epistemic standards. For DeRose, this is because, normally, S' belief that they have hands need only track the truth within a localised sphere of possible worlds that quite closely resemble the actual world, @, and S is, in fact, in a very strong epistemic position with regards to this fact: you have to journey to some very exotic, remote worlds before you find a scenario in which S will falsely believe they have hands, or disbelieve it when they do. Accordingly, when epistemic standards are normal, S's belief they have hands is safe enough for knowledge. Meanwhile, for Cohen, S ordinarily knows they have hands because their evidential corpus is mutually exclusive with every epistemically relevant alternative to that possibility. Being a BIV, for example, *isn't* normally relevant, because it's too improbable, and isn't salient to the conversational participants. Thusly, contextualism "explains our confidence in the truth of our everyday attributions of knowledge" (*ibid.*, p.96), and the intuition that 'scepticism' is false, especially when it's considered in isolation.

On the other hand, however, uttering or considering 'ignorance', just as it can make that claim come out true in your context, can also make 'scepticism' come out true. For DeRose, this is again because, in accordance with the rule of sensitivity, it expands the sphere of epistemically relevant worlds to include the closest worlds where S is a BIV. We can assume that, in a BIV world, S would continue to believe (now falsely) that they had hands. Ergo, they fall short of the new epistemic standards for knowing they have hands: their belief in this does not track the truth over a great enough distance from @. Meanwhile, for Cohen, because uttering 'ignorance' makes S being a BIV a salient, and therefore epistemically relevant proposition for them, it becomes a relevant alternative to them having hands. Again, their evidence doesn't eliminate this alternative, so they fail to know they have hands. Thus, contextualism explains the intuitive pull of AI.

The reason the contextualists give for us feeling confused over 'scepticism' in the face of AI is because of 'semantic blindness'. Although we intuit that the sceptic's conclusion is true, we also

sense that it would come out false if uttered in most, ordinary contexts, and simply fail to recognise the compatibility of these facts, because we fail to notice that the truth-conditions of our K-sentences have shifted. Cohen (2005) summarises the phenomenon thusly:

> "Although 'knows' has a contextualist semantics, competent speakers can be blind to this fact" (p.208).

Thus, contextualists promise to explain our conflicted, fluctuating intuitions about 'scepticism'.

They might also use semantic blindness to explain why, as DeRose says, our intuitions about 'ignorance' can likewise become conflicted. Given their commitment to 'closure', and to S knowing they have hands by ordinary epistemic standards, contextualists are also resigned to saying that S knows they aren't a BIV, again by ordinary standards. Indeed, DeRose (2017) explicitly says that: "on my view, we *do* know that we're not BIVs by the epistemic standards that govern most conversations" (p.189, emphasis in original). For DeRose, for example, this knowledge is achieved because, assuming S isn't a BIV, none of the epistemically relevant worlds are ones in which the BIV scenario obtains, because that possibility holds in only distant, exotic worlds – where much of what would be true in a ~BIV world is, instead, false – and ordinary epistemic standards only require S' beliefs to track the truth across quite close worlds. Accordingly, S's belief that that they aren't a BIV is typically safe enough to qualify for knowledge[22]. For DeRose, this then explains why we sometimes feel that 'ignorance' is false: we sense that, normally, S does, in fact, know that they're not a BIV. The reason we can feel conflicted about it, he says, is again due to semantic blindness: we sense that 'ignorance' is normally false, but true in the context where it's been uttered, and fail to recognise the reconcilability of these facts.

Finally, both contextualists regard their diagnosis of the Harman-Vogel paradox to be genuinely anti-sceptical. They do, of course, allow certain denials of paradigmatic knowledge, such as 'scepticism', to occasionally come out true, when unusually demanding epistemic standards are in play. However, those denials are still false in almost every other situation. As DeRose (2017) explains:

> "[T]he fact that the skeptic can install very high standards that we don't live up to has no tendency to show that we don't satisfy the more relaxed standards that are in place in more ordinary conversations and debates" (p.29).

[22] There's a lot more to DeRose's story about how S comes to know ~BIV, such as how that belief is actually acquired (in an epistemically non-vicious way), and how it's justified. Some of these further details will come up in later sections of this chapter.

And both theorists consider the real threat of scepticism to be its claim that most of our ordinary clams to know things are false. As Cohen (1988) says:

> "What is truly startling about skepticism, is the claim that all along, in our day to day lives, when we have claimed to know things, we have been wrong-we have been expressing propositions that are literally false" (p.117).

As such, because their view says that most of our everyday knowledge claims come out true, it entails that scepticism is false in every important and interesting respect.

To reiterate, the fact that EC provides an anti-sceptical reply to AI – and one that can comprehensively resolve the Harman-Vogel paradox – is meant to be one of the most significant considerations in favour of a contextualist epistemology. In the next section, I'll argue that, when making this case for EC, DeRose and Cohen both practice a centralist methodology, citing intuitions in a crucial evidential capacity.

2.3 Contextualism and Centrality

In this section, I argue that, when making the above case for contextualism, DeRose and Cohen both practice a centralist methodology, citing intuitions in a crucial evidential capacity. In fact, I identify multiple forms of centrality at work in their argumentative method. In later chapters, I'll develop an account of what the relevant intuitions are, defending the view that they're a sub-category of what psychologists call 'epistemic feelings', and thus that said feelings have been highly influential to the evolution of modern epistemology.

2.3.1 The Argument from Intuition-Talk

Despite the lack of any *explicit* defences of centrality (prior to those written in response to his PWI), Cappelen (2012) suggested that there are at least two arguments for it that philosophers might have been *implicitly* relying on for some time. One of these is the 'argument from intuition-talk' (or AIT)[23]. The AIT begins with the observation that contemporary analytic philosophers frequently use intuition-talk, and often to modify propositions, employing, paradigmatically, phrases like '*p* is intuitive' or 'intuitively, *p*'. Then comes the claim that, furthermore, the propositions marked as 'intuitive' are typically used as the premises of an argument, or otherwise

[23] Although this argument had not been made explicitly before its appearance in PWI, versions of it have since been presented by Cappelen's centralist interlocutors. See, for example, Bengson (2014).

presented as considerations in favour of a thesis[24]. Finally, we're told that the best explanation of these facts is that intuitions are playing an evidential role in philosophical arguments: the authors are citing the intuitiveness of their premises/considerations as the evidence or justification for them, and thus, in turn, as the support for their conclusions.

The beginnings of an AIT can be made for 'contextualism centrality' – to wit, the use of intuitions as evidence in the case for EC specifically – by noting just how frequently intuition-talk appears in Cohen and DeRose's work. You'll notice that many of the quotes in the previous sections employ such language, and often to modify propositions, talking about the 'intuitive pull' of AI's premises, the 'intuitions' underwriting the closure principle, or the 'intuitive power' of the sceptic's conclusion, for example. Indeed, it's so difficult to fairly explain their argument for EC without using quotes that feature such terminology, that I finally elected not to do so, even if this can make my presentation of their case seem somewhat question-begging in favour of my centralist interpretation of their method.

However, considering the arguments that have since been presented, by anti-centralists, against an evidential reading of intuition-talk, and given the heterogeneity of centrality, discussed in chapter 1, a lot more needs to be said: we have to be certain that contextualists' intuition-talk really is citing evidential sources, and, if so, in precisely what sense(s) such intuitions are being used 'as evidence', respectively.

In the remainder of this chapter, I analyse the contextualists' use of intuition-talk in much more detail, identifying several, distinct ways in which they employ intuitions as evidence, and reply to both real and anticipated anti-centralist responses to my interpretation.

2.3.2 Explanation Centrality

As I've made clear, the crux of the case for EC is its ability to explain various facts about AI, while avoiding outright scepticism. It's also clear that, for the contextualists, those facts are *intuitions*. However, since this last point is especially important for what follows, I'll say a bit more in support of that interpretation.

Recall that DeRose (2017) describes our ambivalence toward 'ignorance' – a key data point that his contextualism is meant to explain – using intuition-talk, as we saw in this passage:

[24] A similar set of claims are made about the corresponding terms 'counterintuitive' and 'counterintuitively', only the propositions that they modify are allegedly presented as considerations *against* a thesis.

> "Though I could always feel a strong *intuitive pull* toward accepting this premise, I found it far from compelling, and I also felt a significant opposing *intuitive push* toward denying it" (p.55, my emphasis).

Likewise, he says that the sceptic's second premise is "underwritten by 'closure' *intuitions*" (*ibid.*, p.54, Footnote 12, my emphasis), and its appeal is, again, one of the facts that his view promises to account for. Moreover, explaining the lure of 'scepticism' is, yet again, a matter of explaining its intuitiveness, as he reveals here:

> "my focus… has been on explaining why it can seem that we don't know: to explain away the *intuitive appeal* of the skeptic's case" (*ibid.*, p.203, my emphasis).

Indeed, when discussing the overall paradox, or puzzle, that AI presents us with (the 'Harman-Vogel' paradox), DeRose often summarises the total data as a collection of intuitions:

> "I recognize attempts by non-contextualists to explain how some of the *intuitions* involved are simply mistaken as a perfectly fine way in principle to address such conflicts" (*ibid.*, p.90, my emphasis).

We can find similar evidence that Cohen, too, calls the relevant data-points constituting the Harman-Vogel paradox 'intuitions'. For example, when discussing our fluctuating, conflicted sense about whether S being a BIV is an epistemically relevant possibility, he says:

> "[O]ne might object that it begs the question against the skeptic to appeal to anti-skeptical *intuitions* about relevance. And there is a sense in which it does. But of course no more than the skeptic begs the question against the relevant alternatives theorist by appealing to his skeptical *intuitions* about relevance" (1988, p.116-117, my emphasis).

Cohen (2005) likewise describes the appeal of 'closure' in terms of what he at times explicitly calls "the *intuitive* appeal of the deductive closure principle for knowledge" (p.205, my emphasis). Moreover, he describes the lure of 'scepticism' as a 'fact about' our intuitions:

> "[H]ere is *a fact about those intuitions*. When we are in a sceptical frame of mind, it seems to us that we have made a discovery - there are lots of things we think we know… that in fact we do not know" (*ibid.*, p.207, my emphasis).

Indeed, like DeRose, Cohen often summarises the sceptical paradox in general as, specifically, a conflict of intuitions, saying things like "this set of *intuitions* presents us with a puzzle" (1988, p.92, my emphasis). Cohen even says that the whole point of his contextualism is to avoid scepticism while also providing an account of how the paradoxical 'intuitions' arise:

> "The burden of the fallibilist is to resolve these puzzles and paradoxes in a way that preserves the truth of our everyday knowledge attributions. But a satisfying resolution requires an explanation of why the paradox arises - an explanation of why we have the *intuitions* that saddle us with the paradox" (*ibid.*, p.94, my emphasis).

So, for both Cohen and DeRose, the relevant facts about AI, which constitute the Harman-Vogel paradox, and which their contextualism promises to explain, are *intuitions*.

So, contextualists defend their contextualism on the grounds that it can *explain intuitions* (as well as circumvent scepticism). Thus, one way that intuitions are used by contextualists is as 'evidence' in the sense that they're data to be explained by their theory. That is, they practice what I earlier called 'explanation centrality'.

To reiterate, practicing explanation centrality doesn't, in itself, entail taking intuitions to be at all reliable. In other words, saying that S intuiting *p* invites explanation doesn't also say that, if S intuits *p*, then *p* is true, or even likely. Indeed, there's some textual evidence that contextualists regard intuitions as *merely* data to be explained, and not as, say, a guide to what's true. For example, they sometimes seem to think that the only thing wrong with invariantist solutions to AI is the fact that they fail to provide a complete explanation of all of the intuitive data-points, not that they contradict those intuitions and are *ipso facto* dubious because intuitions are a reliable guide to truth. For example, Cohen (2005) rejects Hawthorn's 'Subject Sensitive Invariantism' (or SSI) because it cannot adequately explain our 'sceptical intuitions' in certain cases of third-person knowledge denials:

> "[T]he SSI theorist's semantic blindness thesis cannot explain why we deny knowledge to others when although error possibilities are salient to us, we know that error possibilities are not salient to them. But *in order to explain our sceptical intuitions, that is the very kind of mistake that the SSI theorist needs to explain*" (p.210, my emphasis).

Here, it's the mere lack of explanation for an intuition that allegedly dooms SSI, not the fact that it contradicts one. Indeed, DeRose (2017, p.33) even implies, at one point, that all that's wrong with simply endorsing AI and adopting scepticism is that the sceptic can't explain our intuition that 'scepticism' is false, nor our intuition that 'ignorance' is true (given that, on their theory, we usually mistake ignorance for knowledge). Here, again, the reliability or unreliability of intuition seems entirely incidental: it's not that the sceptic contradicts intuitions that makes their position so implausible, but just the fact that they have no explanation for (some of) them.

So, there's some evidence that contextualists regard intuitions as merely data to be explained, and do not endow them with any greater evidential status. As I'll next show, however, that's an oversimplification of their method.

2.3.3 Intuitive Damage Control

As I'll now demonstrate, both Cohen and DeRose prefer explanations that aren't merely comprehensive, but also 'damage controlling'.

DeRose (2017) acknowledges that it's perfectly acceptable for a philosopher to reject an intuitive proposition, providing one's theory can "explain away why we can find it so plausible" (p.89). In other words, he allows for debunking, rather than ratifying accounts of where our intuitions come from. However, crucially, he seeks explanations that minimise how many intuitions he has to reject, and how strongly he has to disagree with them. In other words, he prefers explanations with the fewest "intuitive costs" (*ibid.*).

We see a similar 'damage controlling' approach in Cohen. Indeed, Cohen (2005) very explicitly states that, out of contextualism and SSI, for example, the *best* explanation for the Harman-Vogel paradox will be the one that does the least violence to intuition:

> "Inevitably the conflict between contextualism and SSI will come down to which view has the greater intuitive costs" (p.207).

The easiest way for an explanation to minimise intuitive costs is to simply agree with our intuitions. One might, for example, say that we intuit *p* precisely because *p*. We see contextualists taking the latter strategy in several places. For instance, when explaining the fact that we typically find 'ignorance' very intuitive, DeRose (2017) says that, on the basis of the 'rule of sensitivity':

> "Since attempts to assert [ignorance] are bound to result in truth, and attempts to deny it are destined to produce falsehood, it's no surprise that we find it so plausible" (p.30).

Here, the explanation that DeRose offers for why we find 'ignorance' intuitive when we hear it is just that it *is* true when we hear it. Likewise, concerning 'scepticism', he says:

> "[W]e seek to explain the persuasiveness of AI, at least in large part, by claiming that the presentation of the skeptic's argument has at least some tendency to put into play the very standards at which we don't count as knowing that we have hands" (*Ibid.*, p.64).

So, again, the sceptic's conclusion strikes us as true precisely because it is.

There are likewise several passages in Cohen where intuitions are straightforwardly explained by the truth of their content. Here, for example, this is especially obvious:

> "[O]ur intuition that the sentence 'Smith knows his car is parked in lot 2' is true *is explained by the fact that at ordinary contexts, the statement is in fact true*. Our apparently conflicting intuition that the sentence 'Smith does not know his car has not been stolen' is false results from our shifting to a stricter context when we consider the possibility that Smith's car has been stolen. At this new context, *the sentence is indeed false*" (2005, p.200, my emphasis).

Similarly, here, our intuition that we know *q* in ordinary contexts is explained by the fact that *h* is not a relevant alternative, such that we do, in fact, know *q*:

> "In everyday contexts where S believes an ordinary proposition q, on the basis of reason r, a radical skeptical alternative h will not be relevant… *This explains our intuition that we know q*" (1988, p.113-4, my emphasis).

So, an easy way to limit the intuitive damage of your explanation is to agree with the intuitions you're explaining. A somewhat subtler way to limit intuitive damage though is to consider even mistaken intuitions to at least provide a rough guide to truth. After all, an explanation that rejects some intuitions altogether is less intuition-friendly than one that treats them as at least approximately correct.

We see contextualists practicing this second strategy as well. Consider, for instance, DeRose's treatment of the sensitivity principle, which states that, if S knows *p*, then S doesn't believe *p* in the closest world(s) where ~*p* (such that sensitive belief is always a condition on knowledge). As we've seen, DeRose considers this to be false. Nonetheless, he acknowledges it to have considerable intuitive support. He says we feel "a strong and surprising intuitive pull toward saying that the subjects… don't know the propositions in question" when their belief is insensitive (2017, p.177), and that, conversely, "modifying the example so that the subject does intuitively seem to know the proposition in question also flips our intuition about the conditional that is crucial to the [sensitivity] account" of knowledge (*Ibid.*, p.138). DeRose doesn't simply reject those intuitions outright, but offers an account of knowledge that makes them at least approximate to the truth. Specifically, he says that holding a safe belief that *p is* necessary for knowing *p*, and, as previously discussed, safety and sensitivity are similar phenomena, and very closely linked. Moreover, he says that certain conversational manoeuvres can, indeed, expand the safety requirement to the point of sensitivity, as per his 'rule of sensitivity'. Thusly, he also provides a charitable explanation for the mistaken intuitions: safety looks a lot like sensitivity, and

sensitivity sometimes is necessary for knowledge, so it's easy to see how one could be confused. DeRose even considers said sensitivity intuitions, though technically false, to still be helpful, precisely because they gesture toward these 'nearby' facts about knowledge:

> "What we learn here will not be limited to some isolated intellectual glitch, but will *point the way to an important lesson about knowledge and knowledge claims* that is of very general application" (*ibid.*, p.63, my emphasis).

So, DeRose chooses to explain even misleading intuitions in a very charitable manner. Thusly, he minimises the intuitive costs of his explanation.

So, contextualists practice a kind of explanation centrality that specifically takes the best explanation for our intuitions to be the one that's most damage controlling, not just the most comprehensive. As I'll next argue, this is tied to them also practicing what I earlier called 'simple centrality'.

2.3.4 Simple Centrality

The contextualist commitment to intuitive damage control is best explained by them taking intuitions to be reliable: explanations that are in closer agreement with intuition are simply more likely to be correct ones, because intuitions are seldom wrong and, when they are, they're often at least approximate to some nearby truth. Of course, this isn't the only possible explanation for their damage-controlling approach, but there's also more direct evidence of contextualists treating intuitions as straightforwardly reliable.

For one, DeRose often cites the intuitiveness of a claim in direct support of it. Recall, for example, that although he doesn't accept the sensitivity principle outright, he does consider it to be approximate to the truth. However, he judges this precisely because said principle is intuitive. He says:

> "Pointing out that someone would have held a certain belief even if it had been false *intuitively* seems a very good reason for denying that they know what they believe" (2017, p.182, my emphasis)[25].

And concludes that:

[25] Elsewhere (*ibid.*, p.182), DeRose compares this to Kripke's claim that "Nozick's third condition [sensitivity] has a clear intuitive basis. "Even if p had been false, you'd still have believed it!" sounds like an objection to a knowledge claim" (Kripke, 2011, p.180).

> "This makes it difficult not to believe that [sensitivity] is at least on the right track" (*ibid.*, p.19).

Conversely, DeRose considers having *counter*intuitive contents or entailments to count against certain theories of knowledge. For instance, when discussing the Nozickian solution to AI, which simply rejects 'closure'-style premises, he says:

> "Nozick's counter-intuitive pronouncement against AI's second premise would count against his theory being victorious" (*ibid*., p.205).

So, for DeRose, intuiting *p* makes *p* (at least) likely, and intuiting ~*p* makes *p* (at least) unlikely.

We can likewise see intuitions being treated as direct support for their content in Cohen. For instance, at one point, Cohen (1986) is discussing whether knowledge requires intersubjectively good reasons, to wit, reasons that lack an intersubjectively obvious defeater. To determine the answer to this question, Cohen constructs a 'test case' in which S possesses a subjectively opaque but intersubjectively evident defeater, *d*, for their reasons, *r*, for believing *q*. Then he says:

> "*Intuitively*, if *S* believes *q* only because obtuseness prevents *S* from discerning the relevance of an obvious piece of defeating evidence, *S* does not thereby know that *q*" (p.577, my emphasis).

The intuition, so identified, is then used to support its own content, namely that subjects with intersubjectively obvious defeaters for their reasons to believe *q* do not know *q*:

> "This suggests that, if *S* possesses an intersubjectively evident defeater of *r* as a reason to believe *q*, then even if the defeater is subjectively opaque, *S* fails to know that *q* on the basis of *r*" (*Ibid.*).

Here, Cohen takes the existence of the intuition that *p* to at least 'suggest', or make likely, *p*. Elsewhere, he relies on intuitions in an even stronger sense. For example, when looking to identify the criteria for epistemic relevance, he says:

> "Whether S knows q will depend on whether any alternatives to q are relevant… Thus the criteria of relevance should reflect our *intuitions* concerning under what conditions S does know q" (1988, p.101-102, my emphasis).

Here, Cohen says that, to identify the relevance of alternatives to *q*, we have to look at where S does and does not know that *q*. However, as we see, for him that just means following our intuitions on the matter of whether S knows. In other words, if we intuit that S knows, then S knows.

So, both Cohen and DeRose treat intuitions as a reliable guide to what's true. That is, they practice what I earlier called 'simple centrality'. To reiterate, I think this best explains why they favour damage-limiting explanations for our intuitions about AI: such explanations are more likely to be accurate, because they're in closer agreement with our, reliable, intuitions.

So, in summation, contextualists don't simply treat intuitions as psychological data to be explained by their theory. Rather, they also treat intuitions as evidence in a more straightforward sense, namely, as a reliable guide to what's true.

We've also seen that they take intuitions themselves ('intuitings'), not just their contents ('intuiteds'), as evidence. On one hand, they seek to explain, not *p*, but why we feel the intuition that *p*, as demonstrated by the fact that they don't always agree with *p*. DeRose, for instance, explains why we intuit the sensitivity principle, not why that principle is true, because he thinks it isn't. On the other hand, they take intuitings to suggest the truth of intuiteds, that is, they don't just take intuiteds to suggest further truths. Cohen, for example, as we saw, takes our intuitions about whether S does or doesn't know *p* as a reliable guide to when S does and doesn't know *p*. Charitably, this isn't just circular reasoning – S knows *p*, therefore S knows *p* – but an inference from (the existence of) the mental state, the intuiting, to its content, the intuited: we intuit *p*, therefore *p*. So, contextualists practice what I earlier called psychological centrality, not just content centrality. Therefore, it certainly isn't mere coincidence that the evidence for EC is intuitive. In other words, contextualists don't just practice what I earlier called 'coincidental content centrality'.

What are the upshots of my findings for the metaphilosophical critique of EC? Well, if the argument for contextualism was just that it could most comprehensively explain our intuitions, then the only relevant evaluative question would be C, namely whether the alleged intuitions actually exist, and thus invite explanation. It wouldn't matter whether they're also a reliable guide to truth, that is, research question D would be superfluous. However, because the contextualist methodology treats such intuitions as reliable, it's equally important that said intuitions prove to actually be so. If they are, then we have reason to favour a damage-controlling explanation like theirs, but, if they're not, then we have much less motivation to accept this over, say, a sceptical resolution to the Harman-Vogel paradox that condemns our intuitions to systematic error. It is, therefore, even more pressing that the exact nature of these intuitions, and their epistemic properties, be determined.

2.3.5 Socratic, Introspective, and Conservative Centrality

Although a damage-controlling explanation centrality, riding on the back of simple centrality, is the most prevalent use of intuitions in the targeted contextualist literature, there's also several other ways in which advocates of EC rely, evidentially, on intuitions. In this section, I acknowledge this yet more complex picture of their philosophical method.

In particular, there are three further forms that centrality takes in the arguments for contextualism, which I call Socratic, introspective, and conservative centrality. These methods are especially evident in the works of DeRose.

2.3.5.1 Socratic Centrality

In this passage, DeRose (2017) speaks of an 'intuitive test' for when one epistemic position is at least as strong as another:

> "[W]e can have… comparative grounds for assenting to conditionals of the form *If S knows that P in A, then S knows that P in B.* In such a case, the comparative grounds for our assent is our realization that S is in *at least as strong* an epistemic position with respect to P in situation B as he is in with respect to that same proposition in situation A, and this comparative conditional serves as a good intuitive test for that comparative fact: It brings that fact to light" (p.23, emphasis in original).

This 'intuitive test' involves considering a 'comparative conditional', which makes explicit, or 'brings to light', a certain 'comparative fact' that was already known, or understood, implicitly. This is very reminiscent of Weatherson's (2013) view that intuitions are 'Socratic knowledge', to wit, a kind of implicit knowledge, or a potential explicit knowledge, that can either be brought to light, or achieved, respectively, by the conscious consideration of certain thought experiments or propositions:

> "[A] subject has *Socratic knowledge* that *p* when they can come to explicitly know *p* by being asked to consider *p* in the right way, perhaps by being told a story that makes *p* vivid" (p.7, emphasis in original).

A potential difference between how Weatherson here understands intuition and how DeRose does in the above excerpt is that Weatherson still seems to regard intuitions as a kind of mental state – being a kind of tacit or potential knowledge state – whereas DeRose, in his talk of 'intuitive tests', seems to regard intuition almost as a kind of faculty, or means, for making certain other beliefs or knowledge conscious (although he might also be happy to call the output of such a process an 'intuition'). Additionally, calling intuition a kind of 'knowledge' suggests that intuitions

are, for Weatherson, factive, and thus infallible (if S knows *p*, then *p* must be true, so, if all intuitions are knowledge, then all have true content), while, as we've seen, DeRose regards at least some of his intuitions to be mistaken. Nonetheless, given the substantial thematic similarities, let's call this incarnation of the practice in DeRose 'Socratic centrality'.

Socratic centrality assigns quite a different role to intuition compared to, say, explanation and simple centrality. Here, intuitions aren't just something to be explained, nor are they straightforwardly treated as a reliable guide to truth. Rather 'intuition' is specifically a means of drawing up to the surface of conscious reflection certain facts that were already known to us.

2.3.5.2 Introspective Centrality

Yet another version of centrality appears in DeRose's discussion of CPK. There, he equates "the correctly refined formulation of closure" with "intuitive closure" (2017, p.163). Initially, this sounds like another, straightforward case of simple centrality, or treating intuitions as reliable: intuitions are trustworthy, so the intuitive formulation of CPK is the correct (or true) one. However, looking closer, something more complicated is going on in this particular case. We can see what when DeRose discusses the infamous 'aggregate epistemic risk' objections to CPK[26]. There, he remarks that formulations of the principle that succumb to such objections don't reveal it to be false, but rather fail to accurately capture it in the first place, precisely because they have these oft-observed 'counterintuitive' consequences. To clarify, for DeRose, the closure principle just is whatever it is that that we all find so intuitive: it's the content of our shared, almost unanimous, pro-CPK intuition. Thus, if a proposed formulation of it has counterintuitive consequences then, *ipso facto*, it's failed to accurately express the principle: it hasn't captured the content of our intuition (and the question of whether the intuited principle is true hasn't even come to the fore). For example, when DeRose rejects what Williamson (2000) calls 'intuitive closure' (p.117), he isn't saying that the proposed principle is untrue. Rather, he just says that calling this 'intuitive closure' is a misnomer precisely because it has infamously counterintuitive implications. In other words, there's an aspect of our closure intuitions that Williamson overlooks (DeRose, 2017, p.174).

Here then, we might say that intuitions are being used as 'evidence' only for *what* their contents are, not for the *accuracy* of said contents. We probe for the contents of our intuitions by, say,

[26] The details of this objection aren't important here. Put very simply, it's to do with the aggregation of epistemic risk across multi-premise deductions. For example, S might be in a strong enough epistemic position to know '*p*' and 'if *p* then *q*', and validly infer *p* from these premises, but still fail to know *q* because the chance of them being wrong about the *conjunction* of '*p*' and 'if *p* then *q*' (that is, being wrong about *at least one* of them) is too high.

reflecting on rival formulations of their content. If a version of a principle is shown, by such means, to have counterintuitive implications, then that's evidence against your view of what the intuition actually is, rather than against the accuracy of the principle you've presented. Let's call this 'introspective centrality', since it's using intuitions to introspect on your own mental contents.

2.3.5.3 Conservative Centrality

The final incarnation of centrality in DeRose's arguments is taking intuitions as justification to believe their contents. For example, DeRose (2017) observes that the way in which we come to believe that we're not a BIV is because of how it 'seems' to us:

> "[T]he way we do ordinarily come to believe [we're not a BIV] is… by our rejecting the BIV hypothesis because it just *seems* so outlandishly (absurdly, ridiculously) false to us" (p.244, my emphasis).

He then asks what it is that justifies that belief, making it a candidate for knowledge. His answer is 'epistemic conservatism':

> "[A]bsent good reason to think things are otherwise, one is justified in believing what seems to them to be the case" (*Ibid.*, p.229).

Here, DeRose states that it seeming to S that *p* (defeasibly) justifies believing *p* for S. Later, in chapter 4, I'll demonstrate that he treats 'seemings' and 'intuitions' interchangeably. Thus, DeRose considers S *intuiting p* to (defeasibly) justify S believing *p*. Let's call this 'conservative centrality'.

To clarify, simple centrality merely says that if one intuits *p* then *p* is likely, but stops short of saying that S merely having an intuition that *p*, absent reasons to think ~*p*, actually *justifies* them in *believing p*. So, while simple centrality presents intuitions as, at best, a good source of evidence to rely on when arguing for a specific theory, conservative centrality grants them the more personal epistemic function of justifying one's beliefs.

So, in summation, the contextualist reliance on intuitions actually takes several, interconnected forms. This is further testament to the aforementioned impracticality of trying to give a more general overview of philosophical method: when we look more closely at even a very narrow band of seminal literature, there's a great deal of variety in how intuitions get used 'evidentially'.

In the last section of this chapter, I'll address several objections to the centralist interpretation of contextualist literature that I've just defended.

2.4 Objections and Replies

In this section, I pre-emptively address five objections that could be made to my centralist interpretation of contextualist literature. The first objection, which I call 'the argument from textual evidence', is an anticipated response that one could make on the basis of certain passages from DeRose (2017) that are harder to reconcile with my view. The other four are existing objections to centralist interpretations of other philosophical literature. The first of those, which I call 'the argument from finite regress', comes from Deutsch (2010). The other three Nado (2015) has labelled 'the argument from unclear application', 'the argument from antipsychologism', and 'the argument from argumentation', and I'll borrow those names for what follows. At the end of this section, I also consider how one might object to the relevance, rather than accuracy, of my interpretation.

2.4.1 The Argument from Textual Evidence

First, I'll address a couple of passages from DeRose's book that superficially conflict with my centralist reading of his methodology.

In some places, DeRose's phrasing seems to caution us *against* interpreting his argument as resting on its ability to explain intuitions. For example, he says:

> "[S]ome may think contextualism is supposed to be primarily supported by its ability to explain how such shifts in intuitions and leanings can occur, or that contextualists themselves seek to support their views by such considerations. And we contextualists have no doubt thought and written things to encourage such an understanding of what we're up to…" (2017, p.91).

Worse still, he elsewhere says:

> "In the end, how most people will weigh the intuitive pushes that they feel against each other and come down on the issue when asked in various ways to vote on it is not all that important—though the evidence reveals that I've always been somewhat curious about that fairly unimportant matter!" (*ibid.*, p.59).

These are likely the passages that any anti-centralist would reach for to support their claim that intuitions aren't, on reflection, that crucial to contextualist arguments and, out of context, they do appear quite damaging, in particular the latter claim that the 'intuitive pushes' DeRose talks about throughout the book aren't 'all that important' after all. Call this 'the argument from textual evidence' against my centralist reading of contextualist literature.

However, we ignore the footnotes at our peril. In this one, DeRose clarifies that, in that last excerpt, what is meant to be unimportant is not the intuitiveness of the sceptic's claims, but the variation in how intuitive those claims are to different people, from different cultures and backgrounds for example:

> "I don't intend this as a general statement about the importance of non-philosophers' opinions about philosophical matters. It is instead a statement about *the relative unimportance of differences in results in this particular situation*, where it seems fairly clear that there are strong intuitive pushes in both directions on a matter…" (*Ibid.*, p.59, Footnote 17, my emphasis).

Here, DeRose suggests that non-philosophers' intuitions *are* important, but clarifies that the variation in just *how* intuitive people find things like 'ignorance', 'closure', and 'scepticism' is incidental, because all DeRose is trying to explain is why there exist 'intuitive pushes in both directions' in the first place, not how the relative strengths of those conflicting pushes are distributed across the population. That is, his contextualism endeavours to explain why our intuitions about sceptical arguments are conflicted *at all*.

We can make similar sense of the other, superficially anti-centralist passage as well. The impression that DeRose is guarding us against is not that contextualists appeal to intuitions, but that they seek to explain, specifically, how 'shifts in intuitions' can occur. Indeed, shortly after that passage, he clarifies that the relevant data is, instead, only that:

> "At least when the claims are considered individually, we do tend to find the claim that we know that O very plausible, but we also can… find each of the skeptic's premises very plausible" (*ibid.*, p.92).

In other words, DeRose is presenting his contextualism as the best explanation for why, especially when considered individually, we can have intuitions both against 'scepticism' and for the sceptic's two premises, and not the specifics of how those intuitions vary with context. His theory explains how semantic confusion gives rise to conflicted feelings about sceptical arguments when they're considered in *any* context.

In general, such seemingly anti-centralist passages should be understood, not as testaments to the irrelevance of intuitions to DeRose's arguments, but as clarifications concerning what it is about intuitions that he specifically seeks to explain. DeRose endeavours to explain the existence of the Harman-Vogel paradox itself, not why the intuitions constituting that paradox are sometimes more or less vivid to certain people, or in certain contexts, for example.

2.4.2 The Argument from Finite Regress

Deutsch (2010) observes that, because research publications are of finite length, their authors won't have sufficient space to defend every claim they make in them all the way down to its epistemic foundations (assuming there are such). Thus, the justification offered for a claim typically 'bottoms out' prematurely. However, Deutsch thinks there are some reasonable end points, where offering further justification indeed seems superfluous, for purely practical reasons. One might, for instance, stop presenting arguments when they reach a claim that's a point of agreement with their interlocutor(s), or that's been sufficiently defended elsewhere in the literature; in either case, offering support for the view would be a misallocation of words. He interprets at least some intuition-talk as simply denoting the end of such an argumentative regress. In other words, Deutsch claims that intuition-talk is used to mark a claim that the author takes to 'require' no further justification in a pragmatic, rather than epistemic sense: it isn't that the claim is foundationally justified, but that defending it is dialectically unnecessary. Because the regress of justification therefore ends with propositions marked as 'intuitive', we might get the impression that their 'intuitiveness' is meant to be the next (and maybe even final) level of justification down, but in fact the language only intends to denote those propositions as a reasonable stopping point, not to cite some consideration in their favour. Ichikawa (2014) also, briefly, suggests something to similar effect, referencing the fact that, if continuously challenged to present justification for their clams, a person might eventually end the regress by exclaiming "I just have an intuition!" (p.4). Here, he likewise interprets the intuition-talk as simply an attempt to halt the train of enquiry. So, for both anti-centralists, intuition-talk isn't really used in an evidence-citing capacity. Call this 'the argument from finite regress' against a centralist interpretation of philosophical literature.

It should be clear why this particular argument doesn't work against a centralist reading of contextualist literature specifically. If contextualists merely called some claims intuitive, and offered no support for them, then yes, potentially their intuition-talk would only be meant to flag claims that they think, for pragmatic reasons, require no defence. However, what they actually do, in several places, as we've seen, is defend intuitive claims, and indeed cite the intuitiveness of those claims, itself, in direct support of them. Recall, for instance, how DeRose explicitly infers that the sensitivity principle is at least a close approximation to the truth from the very fact that it's intuitive, or how Cohen defends certain claims about what's known by appealing to our intuitions about what we know. In both cases, the contextualists' views aren't justified by something else, which they call intuitive, and then offer no support for; rather, the claims are called intuitive, and supported by that very fact. In other words, it's not that 'intuitive' claims are those left without support, but rather the intuitiveness of claims that serves as their support.

Moreover, there's ample evidence that contextualists use intuition-talk to cite mental states. I'll say more about this in the next chapter, but already we can see that, in some of the previously quoted excerpts, they refer to intuitions as conscious phenomena: something that they 'feel'. Thus, they use intuition-talk to speak of a mental event, rather than to simply mark a proposition as the termination of a dialectical regress.

2.4.3 The Argument from Unclear Application

Some philosophers have claimed that intuition-talk is defective, *qua* academic terminology, because it's too vague. Williamson (2007), for example, thinks that, because it's employed in a tremendously generic, unspecific way, intuition-talk serves to 'fudge', rather than elucidate matters:

> ''[P]hilosophers might be better off not using the word 'intuition' and its cognates. Their main current function is not to answer questions about the nature of the evidence but to fudge them, by appearing to provide answers without really doing so'' (p.220).

The claim that intuition-talk is defectively underspecifying is, however, most strongly associated with Cappelen (2012). Cappelen has observed that, even in common parlance, we use words like 'intuitive' to describe a very wide and heterogeneous range of things, from activities, like playing games (I find playing some video games more intuitive than others), to operating systems (I find Windows 10 more intuitive than Windows 8), to people (my friend Giulia is more intuitive than my friend Joe)[27]. This variety of applications, in turn, suggests a corresponding variety of meanings. When I talk about the relative intuitiveness of playing different games, I seem to mean how unconsciously or 'instinctively' I can do something. When I describe Windows 10 as intuitive, I seem to mean that everything is, and works, as I expect: it doesn't subvert my expectations. When I describe people as intuitive, I seem to be referencing their emotional insightfulness or empathic ability.

Now, it could be argued that, despite its ambiguity in ordinary English, intuition-talk has acquired a single, consistent meaning among analytic philosophers. However, Cappelen denies this. On the contrary, he claims, the meaning of words like 'intuition' displays just as much ambiguity in 'philosophers' English' as it does in common parlance, if not more so. Indeed, as I explained in Chapter 1, this much can be hard to deny, due to the massive degree of variation in how philosophers' intuition-talk can be, and has been, interpreted.

[27] See PWI, Chapter 2 for Cappelen's own examples.

According to Cappelen, owing to its massive ambiguity, unless a particular meaning is explicitly stipulated, the use of intuition-talk in academic contexts is, again, 'defective': a hopelessly imprecise terminology. Indeed, with no standardised meaning to default to, such unreflective uses may even amount to nonsense: "there's no semantic anchor point and the term ['intuition' or 'intuitively'] fails to have a semantic value" (*Ibid.*, p.50).

Cappelen turns this into a case for anti-centralism. Professional analytic philosophers, the argument goes, are intelligent, conscientious, methodical individuals, who surely wouldn't rely on such defective language in defence of a thesis. Moreover, analytic philosophers are especially preoccupied with clarity, and so expected to provide concise definitions of their key terminology. Thus, if they meant to use intuition-talk to do anything substantial – if it was intended to perform any significant argumentative work – then they would, at least, offer a clear definition of it, to compensate for its inherent ambiguity. Otherwise, we should assume that the real work is being done by something different, something much more concrete than this hazy, undefined thing called 'intuition'. Given that most analytic philosophers have no such proclivity to offer explicit definitions of their intuition-talk, the most charitable interpretation is, therefore, that such language is incidental to their arguments. So, charitably, intuition-talk isn't citing evidence[28]. Call this 'the argument from unclear application' against a centralist interpretation of philosophical literature.

In response to this argument, we can simply point to the overwhelming textual indications, already presented, that contextualists do use intuition-talk to cite evidence. If that language transpires to be defectively ambiguous or underspecifying, this won't be sufficient to defeat such an interpretation, and will merely reflect badly upon their argument. After all, if contextualists' intuition-talk is defective, then their argument lacks clarity or even coherence. I appreciate the appeal to interpretive charity implicit in 'the argument from unclear application', but surely such a principle is limited: we shouldn't embrace an implausible interpretation of the literature, with no textual evidence whatsoever supporting it, just to avoid attributing incomprehensible arguments to philosophers, nor forget that, sometimes, arguments really do rely on defective terminology.

[28] Cappelen's alternative is that intuition-talk is purely stylistic, adding nothing to the content of an argument. According to this, intuition-talk is as something that can, in principle, be removed from our texts altogether, without changing the content or structure of the arguments, only their presentation. Cohnitz (2012) summarises Cappelen's proposal as a "charitable reinterpretation of that defective discourse" (p.4).

2.4.4 The Argument from Antipsychologism

Some anti-centralists have argued that philosophers don't cite intuitions, or facts about them, as the evidence for their views, because they don't rely evidentially on mental states at all, but on something else, something non-psychological. Call this 'the argument from antipsychologism'.

There are multiple versions of this argument, depending on what one thinks philosophers make evidential use of instead of mental states. Cappelen (2012), for example, thinks that philosophers often appeal to facts about language. He points us to Burge's (1979) arguments for content externalism, for instance, which are widely regarded to involve appeals to intuition, and observes that, when discussing his 'arthritis' thought experiments, Burge only relies on a claim about how people normally speak, or what they would typically say in a given context – he "makes an initial set of observations about usage… that is an *empirical fact about speech behavior*" (Cappelen, 2012, p.143, emphasis in original). Thus, Burge's argument appeals to ordinary language, not to a mental state.

We can locate another version of this argument in Williamson's *The Philosophy of Philosophy* (2007). Describing, primarily, the 'method of cases' – roughly, presenting hypothetical scenarios and inviting the reader to form, or concur with, certain judgements about them – Williamson (2011) claims that, when forming judgements about such cases, we don't employ "a mysterious *sui generis* faculty of rational intuition, or anything of the kind" but rather "more mundane skills, such as careful attention to details in the description of the scenario and their potential relevance to the questions at issue" (p.217). When presented with a Gettier case, for example, we might first consider its logical possibility, simply by checking for contradictions in the case description, which could also involve employing a cultivated sensitivity to the fine details of the vignette. Then, we might compare it to other, previously encountered cases and apply some inductive reasoning: is this like other cases of knowledge? Finally, we evaluate whether our narrow conclusions about the scenario really do entail the broader conclusions of Gettier's argument against the 'justified true belief' theory of knowledge, and that just means employing our understanding of validity. By stressing the application of ordinary reasoning faculties, over the appeal to some peculiar mental phenomenon called 'intuition', or its contents, Williamson is likewise arguing that we don't 'psychologise' our evidence.

The first problem with this style of argument is that it presupposes what philosophers mean by 'intuitions' – they're not facts about ordinary speech behaviour, for example, but a mental state, and, according to Williamson, a 'mysterious' one at that. But, so far, I've been neutral about what contextualists take the 'intuitions' they talk about to actually be. At this point, they may, indeed, just be facts about ordinary language, for example. What matters to me, in this chapter, is just

whether they're being cited in an evidential capacity. To clarify, there's two ways that one can approach research question A (is centrality true, or, do philosophers appeal evidentially to intuitions?). On the one hand, we can ask whether certain philosophers evidentially appeal to what many theorists already take intuitions to be, whether those certain philosophers personally call them 'intuitions' or not. On the other hand, we can ask whether certain philosophers evidentially appeal to something that they themselves often call 'intuitions', whether that 'something' is what others think of as intuitions or not. Advocates of 'the argument from antipsychologism' seem to be taking the former approach, trying to score a spiritual victory over centralists, by demonstrating that philosophers don't appeal to what many theorists, including most centralists, take intuitions to be. However, my project takes the second angle: having observed the growing prevalence of intuition-talk in contextualist literature, my goal is to unscramble what this language is doing – is it citing evidence and, if so, what is that evidence? Thus far, I've merely answered the former question, in the affirmative.

Moreover, it's clear enough that 'the argument from antipsychologism' doesn't work against a centralist reading of contextualist literature specifically. To reiterate what I said in section 2.4.2, Cohen and DeRose do often describe intuitions as mental states, referring to them as conscious phenomena that they 'feel', and, as we've seen, they cite that feeling, directly, in support of their claims. Thus, they do in fact appear to 'psychologise' their evidence. Again, I'll defend this claim more comprehensively in the next chapter.

2.4.5 The Argument from Argumentation

It's quite common to claim that intuitions can provide 'foundational justification': that they can justify further cognitive states (such as beliefs based on those intuitions), but do not, themselves, require justification from anything else, 'lower down'. Here, for instance, is BonJour (1985), defining 'intuition' as something with precisely this feature, which he calls 'givenness':

> "The basic idea of givenness... is to distinguish two aspects of ordinary cognitive states, their capacity to justify other cognitive states and their own need for justification, and then to try to find a kind of state which possesses only the former aspect and not the latter—a state of immediate apprehension or *intuition*" (p.78, my emphasis)[29].

If we accept that intuitions are foundationally justified, then, if philosophers were using the intuitiveness of *p* as their evidence for *p*, they'd have no need to provide any additional justification for *p*: their argument would 'bottom out' at their intuitions, since they require no

[29] Bonjour goes on to deny that givenness – and hence 'intuition' – is possible.

further support. According to some anti-centralists, however, philosophers do give arguments for allegedly intuitive claims and, therefore, don't really support them on the basis of their intuitiveness. Call this the 'argument from argumentation' against a centralist interpretation of philosophical literature.

Cappelen (2012), for example, points us to Lehrer's famous thought experiment, in which a person, Mr. Truetemp, has, unbeknownst to him, had a device implanted in his brain that allows him to tell the temperature in the room with perfect reliability. The details of Lehrer's argument aren't relevant here, but it's commonly assumed that a certain judgement about this vignette – that Mr. Truetemp doesn't *know* the temperature – is meant to be intuitive, and that Lehrer supports said judgement on the basis of its intuitiveness. However, Cappelen observes, Lehrer actually provides an argument for the claim that Mr. Truetemp doesn't know, in the form of this passage:

> ''More than possession of correct information is required for knowledge. One must have some way of knowing that the information is correct'' (Lehrer, 2000, p.188).

Again, if the claim was meant to be supported by its intuitiveness, then there wouldn't be any need to provide this further justification for it. Therefore, the support for Lehrer's claim isn't, on reflection, intuition. Or so goes Cappelen's argument.

Ichikawa (2014, c.f. Ichikawa and Jarvis (2013 esp. ch. 12-13)) mounts a similar attack on centrality. Asking us to consider what a philosopher would actually say in defence of the claim that the subject of a Gettier case, S, doesn't know some proposition, *p*, for example, he argues that they wouldn't, in fact, cite the intuitiveness of the claim, but rather, "for instance, the fact that S's belief that *p* was derived from a falsehood, or that he was lucky to have gotten his belief right" (Ichikawa, 2014, p.4). Again, this is meant to show that philosophers support allegedly intuitive claims with further arguments, rather than relying on the foundational justification given to them by their intuitiveness.

To clarify, advocates of 'the argument from argumentation' don't assume that we can't support the same claim using both arguments and intuitions. A proposition can receive justification from multiple, different sources, so philosophers might intend such arguments to lend the intuitive justification an extra boost, or *vice-versa*. This possibility is explicitly acknowledged by Cappelen throughout PWI, and he makes a point of reiterating it in later work:

> "I tried hard… to make sure readers wouldn't think that I assumed that if an argument is given for *p*, then *p* is not supported by an intuition" (2015, p.29).

Nonetheless, Cappelen takes it as a 'rough guide' to the reliance on intuitions in an argument if no further case for an allegedly intuitive premise is provided. Conversely, if *p* is supported by an argument, then we have little reason to suppose that it's also supported, foundationally, by its intuitiveness.

As with 'the argument from antipsychologism', one issue with 'the argument from argumentation' is that it presupposes what philosophers mean by 'intuition'. Again, at this stage, what contextualists mean by 'intuition' may or may not be something that provides, or which they take to provide, foundational justification. All that matters for my purposes here is whether they rely on whatever they call 'intuition' in an evidential capacity.

Of course, the bigger problem with 'the argument from argumentation' is that contextualists don't, in fact, provide arguments for the claims they call intuitive. Again, this is especially clear where DeRose says that the sensitivity principle is intuitive, and then takes it as an approximate truth without further argumentation. Indeed, as we've already seen, DeRose infers directly from the intuitiveness of the principle to its likelihood. Similarly, Cohen, to reiterate, infers what we know or don't know directly from our intuitions about such matters, and again provides no further support for such claims. Thus, these cases satisfy Cappelen's 'rough guide'.

This also serves to highlight just how important intuitions are to the contextualists' argument. If they cited intuitions alongside further considerations, then the former might be playing a merely supplementary role. As it happens, they are, in most places, the sole support for contextualists' claims.

2.4.6 Scope and Relevance

Some philosophers – such as Sosa (2000) – have claimed that contextualism isn't really an epistemological thesis after all, but rather something that belongs to the philosophy of language, or even linguistics, because it's a claim about the semantics of K-words, not about the metaphysics of knowledge itself. If one has sympathy for this view, then they might question the scope and relevance of my own metaphilosophical research. Specifically, the objection could be made that, because contextualism isn't representative of classic epistemology (or even philosophy), my analysis of contextualist arguments in this chapter actually reveals very little, if anything, about the versions of centrality practiced across the discipline as a whole, and the interpretation of contextualists' intuition-talk that I provide later in the thesis likewise won't reflect how that language is used by epistemologists more broadly.

I have two things to say in response to this anticipated worry. First, while the claim that K-words are semantically context-sensitive is linguistic, there are many other claims making up the contextualist position that are paradigmatically epistemic. Contextualists combine their semantic thesis, about the context-sensitivity of K-words, with other claims, about knowledge, to provide an anti-sceptical solution to the Harman-Vogel paradox. This strategy is made clear in this passage, for example, where DeRose (2017) presents us with what he takes to be the guiding question of his recent book:

> "Can the contextualist semantics of 'know(s)' being proposed… *together with [an] account of what knowledge is*, well explain, and explain better than rival accounts, how the premises of the particular skeptical argument being addressed can be as plausible as they are, while its conclusion is as implausible as it is?" (p.92, my emphasis).

So, the contextualists' explanation combines both semantic and epistemic claims: a claim about the 'semantics of 'know(s)'' (that K-words are context-sensitive) and claims about 'what knowledge is' (that it requires safe belief, for example). Therefore, it's oversimplified to regard the contextualist position as a purely linguistic one.

Moreover, if anything, the semantic claim, that K-words are context-sensitive, isn't intuitive. Throughout the history of analytic philosophy, K-words have seemed, even to trained epistemologists, to be semantically invariant, not context-sensitive. Indeed, the latter idea is so hard for most people to swallow that Schiffer (1996) and Hofweber (1999), among others, have made this into a criticism. We would, they claim, sense, or intuit, that K-words were context-sensitive if they were, and we don't – rather, they seem to be exactly the opposite – so EC is implausible. Thus, it's highly improbable that contextualists take the intuitiveness of the semantic claim, specifically, to lend any credence to their wider view, since it isn't intuitive to begin with[30]. By contrast, the intuitions that, as we've seen, *are* appealed to in 'the argument from anti-scepticism' are invariably about epistemic matters. Our intuitions about 'ignorance' and 'scepticism' are straightforwardly about what S knows, for example, while out intuitions toward CPK, safety, and sensitivity are to do with epistemic principles. So, it's not linguistic intuitions that support EC here, but epistemic ones.

[30] This isn't to assume that some of the claims they elsewhere use to support the semantic thesis aren't meant to be intuitive, and those claims (such as the ones made in DeRose (2009)) are more plausibly of a linguistic nature – claims about, for example, whether certain utterances are 'good English' when uttered in certain contexts. However, when mounting 'the argument from anti-scepticism', specifically, contextualists offer no such linguistic defence of the context-sensitivity of K-words.

So, my first reply to this anticipated worry is that, although the contextualist treatment of the Harman-Vogel paradox (crucially) relies on a semantic claim, their position incorporates epistemic claims as well, and the intuitions that I'm focussing on are about the epistemic matters, not the linguistic ones. There's no reason, therefore, not to treat them as representative of intuitions in epistemology.

My second, briefer reply is that, even if we didn't regard the intuitions I'm looking at to be representative examples of epistemic intuitions, those intuitions still support an argument that, as I explained in Chapter 1, has been extremely influential to contemporary epistemology. Therefore, even if these intuitions weren't representative, looking at how they're used, and what they are, would still teach us about how intuitions shape modern philosophy.

2.5 Chapter 2 Conclusion

In conclusion, advocates of epistemic contextualism employ intuitions in an evidential capacity, and indeed in several different ways: intuitions are used as data points to be explained, as a guide to truth, as justification to believe their content, and so on. This, then, is my partial answer to research question A: intuitions indeed play an evidential role in our discipline. To reiterate what I said in Chapter 1, while this isn't enough to show that centrality, as a generalisation about the field of philosophy as a whole, is true, the fact that these particular intuitions have been used as the primary support for contextualism, which is having a burgeoning influence on modern epistemology, and that epistemology is of ubiquitous philosophical significance, with ties to arguably all other core philosophical topics, such as ethics and metaphysics, means that treating intuitions as evidence has profoundly shaped the discipline at large, contrary to the views of many anti-centralists. In other words, as I stated in Chapter 1, the 'spirit' of centrality, at least, is true: intuitions *have* been 'central' to the history and evolution of modern philosophy, and so the study of said intuitions *could* have profound metaphilosophical implications.

So, philosophers use intuitions as evidence, and doing so has had a significant impact on modern philosophy. In the next chapters, I'll develop a novel account of what these intuitions are. I'll begin, in Chapter 3, by arguing against 'minimalist' conceptions of philosophical intuitions. Then, in Chapter 4, I'll defend the view that intuitions are a variety of 'intellectual seemings'. In Chapter 5, I'll argue that intellectual seemings are not *sui generis*, as many are wont to assume, but rather a sub-category of metacognitive experiences that psychologists call 'epistemic feelings'. I'll conclude that such feelings have, therefore, significantly shaped modern philosophy and, in Chapter 6, consider some of the implications of this discovery.

Chapter 3 Against Minimalism

As I established in Chapter 2, when making 'the argument from anti-scepticism' for epistemic contextualism (or EC), contextualists rely evidentially on intuitions, in a number of different ways. Since this argument is a primary motivation for EC, and EC has had a significant impact on contemporary epistemology, this centralist method has played, and by all accounts continues to play, a significant role in shaping our discipline. If we want to accurately evaluate that influential methodology, however, we first need to clarify what the 'intuitions' in question actually are. Only then can we confidently assess whether people really have these alleged 'intuitions', and whether they're a good source of evidence for philosophical theories or not. Accordingly, the remainder of this thesis is devoted to an interpretation of contextualists' 'intuition-talk', and aims to unravel exactly what is meant, in *this* literature, by words like 'intuition' and 'intuitive'.

This won't be an easy task, however. As we'll see, contextualists describe intuitions in many different ways. The relevant phenomena are, of course 'intuitions', but they're also 'beliefs', 'judgements', 'verdicts', 'attributions', 'reactions', 'pushes', 'pulls', 'impulses', 'inclinations', what we 'say', how things 'seem', and more. Meanwhile, a proposition, *p*, being 'intuitive', or having the property of 'intuitiveness', is also described as *p* being 'plausible', 'persuasive', 'compelling', 'difficult to deny', and so on. Moreover, it's hard to see how one might reconcile all these different ways of speaking. Some, such as belief-talk, judgement-talk, and verdict-talk, seem to suggest that intuitions are just a kind of belief. Others, though, such as inclination-talk, suggest that having an intuition is less epistemically committal than that, since being inclined to believe something need not entail actually believing it. Others again suggest that intuitions have more to do with subjective confidences, or credences: how 'plausible', 'persuasive' or 'compelling' we find something. Furthermore, beliefs, inclinations, and credences are all (typically) unconscious states, while other descriptions make intuitions sound more distinctly phenomenal, such as the talk of them as 'impulses' and 'feelings'. Also, interchanging intuitions with 'what we would say' can make them sound like speech acts, rather than mental states: nothing more than empirical facts about ordinary speech behaviour. And so on. As we'll also see, it's not even that a particular 'intuitive data point' is consistently referred to using one set of these terms: the very same intuition will, at times, be spoken of in one way and, at other times, in quite a different sounding way. Indeed, the different ways of describing intuitions are interchanged so casually – with authors fluctuating between them in the same paragraph, and sometimes even the same sentence – that no good case can be made for the interpretation that there are just multiple, distinct things being referred to here.

What we're confronted with, then, is a considerable degree of inconsistency and imprecision in how contextualists describe intuitions. In light of this, I'll push for a 'rational reconstruction' of their terminology. To wit, I'll argue for a way to cash out their intuition-talk in clearer, more precise terms, and reveal the phenomenon that they've been vaguely circling around with their various, sometimes conflicting descriptions of intuition.

More specifically, over the next three chapters, I argue that the relevant intuitions are best understood as 'intellectual seemings', and that these are, in turn, not *sui generis*, but a particular sub-category of what psychologists call 'epistemic feelings'. Crucially, my claim is not that this is what analytic philosophers in general use intuition-talk to mean, as different philosophers across the discipline might employ that language in a multitude of different ways, as previously discussed. Rather, my claim is only that this is what contextualists, roughly, have in mind, whether they realise it or not. How representative this phenomenon is of philosophical intuitions more broadly is a matter for future debate, but, again, given how influential these particular intuitions have been, it'll be sufficient to understand their nature for now.

I begin, in this chapter, by arguing against 'minimalism about intuitions' (herein, just 'minimalism'). Minimalism refers to the view that intuitions aren't a *sui generis* mental state, but reducible to another, already widely recognised state, such as beliefs. The position is minimalist in the sense that it's metaphysically parsimonious. To clarify, while my own, novel account of intuitions ultimately agrees that they aren't *sui generis* – being reducible to epistemic feelings – I'll use this space to argue against extant minimalist theories.

Keeping in mind that I'm investigating the intuitions being evidentially relied upon in 'the argument from anti-scepticism' for EC, I argue, specifically, against the views that said intuitions are reducible to outright beliefs (§3.1.1), partial beliefs (§3.1.2), dispositions to believe (§3.2), or facts about ordinary language (§3.3). While these certainly aren't the only minimalist interpretations of intuition-talk that have been offered – another example from the literature is the claim that intuitions are items in the 'common ground' between authors and their interlocutors (Cappelen, 2012) – they are each, as I'll show, readings with some *prima facie* plausibility, having at least a modicum of textual evidence from the seminal contextualist literature in their favour, and thus warrant serious consideration.

Later, in Chapter 4, I'll argue that the relevant intuitions are, instead, intellectual seemings. Then, in Chapter 5, I'll make the case that said seemings are a sub-category of epistemic feelings. Finally, in Chapter 6, I'll consider some of the implications of this discovery.

3.1 Against Doxasticism

Doxasticism about intuitions (herein, just 'doxasticism') is the view that intuitions are just a kind of belief. This view is motivated primarily by its metaphysical parsimony, but also by the apparent correlation between what we believe and what we intuit. To elaborate, take, for instance, the set of propositions i-v:

i) If ~~*p*, then *p*
ii) Torturing a sentient being for fun is immoral
iii) A square cannot have five sides
iv) 3 is more than 2
v) Identity is transitive (if A=B, and B=C, then A=C)

i-v are drawn from Pust (2017) and Bengson (2015), who list them as paradigmatic examples of intuitive propositions. i-v are also propositions that, we can assume, almost everyone believes (at least after they've considered them). Indeed, whenever a subject, S, intuits a proposition, *p*, we would normally expect S to believe *p*, and this very correlation is sometimes thought to suggest doxasticism. As Koksvik (2012) puts it, "some take the obvious fact that there is a fairly tight connection between intuition and belief as incentive to say that intuition *is* a belief" (p.12, emphasis in original).

In this section, I argue against reading contextualists' intuition-talk, specifically, as referring to beliefs, whether outright (§3.1.1) or partial (§3.1.2) ones.

3.1.1 Against Strong Doxasticism

What I'll herein call 'strong doxasticism' is the view that intuitions are 'outright' beliefs, rather than partial beliefs (or degrees of confidence). This position is sometimes attributed to Lewis (by Pust, 2017, for example), on the basis of this passage, wherein he declares that our intuitions are 'opinions':

> "Our 'intuitions' are simply opinions… Some are commonsensical, some are sophisticated; some are particular, some general; some are more firmly held, some less. But they are all opinions…." (Lewis, 1983, p.X).

The strongest version of strong doxasticism would be to say that beliefs and intuitions are type-identical, such that all beliefs are intuitions, and *vice versa*. Call this 'type strong doxasticism'. However, it's generally accepted that people can believe unintuitive things, and even things that are neither particularly intuitive nor unintuitive. Therefore, not every belief that *p* has a

corresponding intuition that *p*, and, if it's possible to believe *p* without intuiting *p*, it can't be the case that *every* belief is an intuition. Thus, type strong doxasticism succumbs to a fairly obvious 'overgeneration objection': it defines intuitions too liberally, and consequently posits the existence of intuitions where there are none. To circumvent this worry, advocates of strong doxasticism typically define intuitions, more narrowly, as a proper sub-set of our beliefs, rather than as beliefs *simpliciter*, such that every intuition is a token belief, but not all token beliefs are intuitions. Call this 'token strong doxasticism'. Token strong doxasticism, unlike type strong doxasticism, is consistent with believing non-intuitive propositions, since only some beliefs – those of the relevant sub-category – are intuitions.

Weinberg *et al.* (2001), for example, define intuitions as, specifically, spontaneous beliefs, to wit, beliefs formed very quickly, before any conscious reflection on their subject matter has occurred:

> "[A]n epistemic intuition is simply a spontaneous judgment about the epistemic properties of some specific case" (p.447).

Likewise, Nichols *et al*. (2003) define an intuition as "a spontaneous judgment about the truth or falsity of a proposition" (Footnote 2). Similarly, Cappelen (2012) suggests, as one of his several, deflationary readings of intuition-talk, identifying intuitions with what he sometimes calls 'snap judgements'.

An alternative is to define intuitions as, specifically, common-sense beliefs. Cappelen (*ibid*., p.71-75), for example, argues that Kripke used 'intuitive' to denote the common-sense views that he held prior to being inculcated with the prevailing philosophical standpoints of his era. The support for this interpretation comes from the fact that Kripke's work often seems to involve contrasting various philosophical views against his pre-informed opinion and, in such cases, the latter are often referred to as 'intuitive' thoughts or ideas. For example, he says in the preface to *Naming and Necessity* (hereafter NN):

> "Our intuitive idea of naming suggests that names are rigid, but… I vaguely supposed, influenced by prevailing presuppositions… ordinary names must not be rigid" (1980, p.4).

Here, Kripke contrasts the view inherited from the 'prevailing presuppositions' in philosophy (which we can therefore think of as the academically informed view), against the idea that held sway beforehand (which we can therefore think of as the pre-informed or 'common-sense' view), and refers to the latter as the 'intuitive' idea. Cappelen (2012) speculates that "a plausible case can be made that Kripke's extensive use of 'intuition'-terminology in [NN] influenced a generation

of philosophers" (p.71) to use it in roughly the same way, and that denoting common sense beliefs is, therefore, a likely interpretation of philosophers' intuition-talk.

Brogaard (2013a) likewise argues that denoting something as a common-sense opinion is one (though not the only, or even dominant) use of intuition-talk in the academy, as in sayings like "intuitively, the speed of light is not a constant" (p.212). Parsons (1995, p.59), meanwhile, simply equates intuitions with items of common sense.

Despite these attempts to weaken strong doxasticism and circumvent the aforementioned overgeneration objection, however, the view that intuitions are outright beliefs remains very unpopular, and has been argued against numerous times in the extant literature (Bealer, 2008; Chudnoff, 2011; Earlenbaugh and Molyneux, 2009; Huemer, 2007; Kagan, 1989; Katz, 1981; Plantinga, 1993; Pollock, 1974; Pust, 2000; Sosa, 2007b; Williamson, 2007; Cappelen, 2015). Nonetheless, there does seem to be some textual evidence that contextualists, at least, use intuition-talk to refer to outright beliefs, and thus the possibility is worth taking seriously.

Next, I'll outline some of that evidence. Then, I'll make four arguments against reading contextualists' intuition-talk as referring to outright beliefs.

3.1.1.1 Textual Evidence of Strong Doxasticism

In this excerpt, DeRose (2017) uses intuition-talk interchangeably with belief-talk:

> "[I]f the skeptic is marshalling deeply felt *intuitions* of ours in an attempt to give us good reasons for accepting his skepticism, it's legitimate to point out that other of our *beliefs* militate against his position…. And if we can further show that those *beliefs* that seem to favor his solution can be accommodated in our solution better than he can accommodate those of our *beliefs* that are hostile to him…" (p.38, my emphasis).

Here, DeRose highlights a conflict between the 'intuitions of ours' that support scepticism and 'other of our beliefs' that contradict it, strongly suggesting that intuitions are just beliefs. Indeed, he casually switches from saying that the support for the sceptic comes from 'intuitions', at the beginning, to saying it comes from 'beliefs', in the closing line.

DeRose (2017) likewise uses intuition-talk interchangeably with judgement-talk. For example, we've already seen that he takes the primary support for the sensitivity principle to be its intuitiveness, but he also, often, states that it's supported by what we 'judge', as he does here:

> "[W]e tend to *judge* that S doesn't know that P when we think that S would believe that P even if P were false…" (p.137, my emphasis).

So, the support for the sensitivity principle is intuition, or, what we judge. Also, the role that said principle is meant to play in explaining the 'intuitive data' constituting the Harman-Vogel paradox (via his Rule of Sensitivity) is sometimes framed in terms of its ability to explain certain *judgements*. Indeed, in this passage, DeRose casually vacillates between the words 'judgments' and 'intuitions':

> "[sensitivity] outperforms its rivals in terms of explaining our *judgments* about what is and isn't known. If I succeed in showing that [sensitivity] is the best explanation for why we have the particular *intuitions* we have, that should motivate us to seek an account of knowledge that makes sense of [sensitivity]" (*Ibid.*).

So, the sensitivity principle explains our intuitions, or, our judgements.

The most natural way to read judgement-talk is as a kind of belief-talk. In particular, it's thought that 'judging' means forming a belief, and 'judgements' are, therefore, the formations of beliefs[31]. For example, Carruthers (2010) says:

> "I take judgments to be events of belief-formation… Judgments are a kind of active, occurrent, mental event, which when stored give rise to dormant, standing-state, beliefs" (p.78).

Thus, when people talk about judgements, they're likely talking, in some sense, about beliefs. So, DeRose's vacillation between judgement-talk and intuition-talk is further evidence of strong doxasticism. It suggests that, for example, the 'intuitions' supporting the sensitivity principle, and the 'intuitions' constituting the Harman-Vogel paradox, are just beliefs.

Meanwhile, in Cohen (1988), we don't find intuitions described as beliefs, and only a single instance where they're referred to as 'judgements', that being here:

> "The non-relevance of skeptical alternatives is a datum provided by our *intuitive judgments* concerning what we know" (p.116, my emphasis).

However, Cohen does sometimes talk about intuitions in terms of what we 'attribute'. For instance, recall that Cohen's strategy for determining the criteria for epistemic relevance is to track what he calls our 'intuitions' about when S knows: if, intuitively, S knows *p* in some context, *C*, but not the negation of some alternative, *A*, then *A* isn't epistemically relevant in *C*, for

[31] Consequently, on the standard reading of judgement-talk, each judgement entails a corresponding belief. Thus, in this thesis, I'll sometimes use/treat 'judgement' and 'belief' interchangeably.

example. However, when he pursues this strategy, what he actually appeals to is, more specifically, the fact that we 'attribute' knowledge:

> "Since ordinarily we *attribute* knowledge in these cases, these alternatives are not relevant in those contexts" (*ibid.*, my emphasis).

So, our 'intuitions' about when S knows are really just our 'attributions' of knowledge. Indeed, when summarising his overall argument for EC, Cohen says that his contextualism can explain a 'pattern of attributions', while preserving closure and avoiding scepticism:

> "What I am proposing is that we can *explain this pattern of attributions* in terms of the mechanism of context-sensitivity in a way that endorses our original knowledge attributions, preserves the truth of the closure principles, and accounts for the power of skeptical arguments" (*Ibid.*, my emphasis).

Given that, elsewhere, Cohen has been very clear that the case for his contextualism is its ability to explain our *intuitions* about AI-style arguments, again in a way that preserves closure and avoids scepticism, we can see that, here, attribution-talk is just replacing intuition-talk.

One way to read Cohen's attribution-talk would be doxastically: to say that, in some cases, we do attribute knowledge to S, and in others we don't, just means that, in some cases, we believe that S knows and, in others, we don't, for example. So, his fluctuation between these terms again suggests strong doxasticism. It suggests that, for instance, the 'intuitions' about when S knows, which reveal the criteria for epistemic relevance, are just beliefs about when S knows, and the 'intuitions' constituting the Harman-Vogel paradox are just a pattern of beliefs.

In the proceeding sub-sections, I'll argue that this textual evidence is misleading, and contextualists' intuition-talk doesn't refer to outright beliefs.

3.1.1.2 The Argument from Lack of Belief

The most common (and straightforward) argument against strong doxasticism is that philosophers don't always believe what they intuit. An oft-cited example of an intuition that's isn't believed by its subject is the intuition that the naïve comprehension axiom of set theory – roughly, the thesis that any sufficiently defined property can be used to determine a set – is true. Many philosophers admit to feeling this intuition, but take it to be mistaken, and resist the urge to believe it, because they're familiar with Russell's paradox (Williamson, 2007, p.217; Cappelen, 2015, p.17; Huemer, 2007, p.31). Bealer (1998), for example, says:

> "I have an intuition – it still *seems* to me – that the naïve comprehension axiom of set theory is true; this is so despite the fact that I do not believe that it is true (because I know of the set-theoretical paradoxes)" (p.208, emphasis in original).

If it's possible to intuit *p* without believing *p*, then the intuition that *p* isn't, itself, a belief. This is an 'undergeneration' objection to strong doxasticism: defining intuitions as outright beliefs fails to capture at least some of our intuitions, because it overlooks those that aren't believed. This criticism is especially popular, in part, because, in contrast to the aforementioned overgeneration objection, it works against both type and token doxasticism: if it's possible to intuit *p* without believing *p* at all, then intuitions aren't identical to *any* token beliefs, let alone beliefs in general. Call this 'the argument from lack of belief' against strong doxasticism.

The same line of argument can be made against reading contextualists' intuition-talk, specifically, as referring to outright beliefs. After all, there are several examples of contextualists failing to outright believe the things they intuit. Recall, for instance, that DeRose finds the sensitivity principle intuitive, but doesn't believe it; in fact, he explicitly rejects sensitivity as a general condition on knowledge. Likewise, consider that, according to the contextualists, when we consider AI, 'scepticism' seems both intuitively true and intuitively false. They claim that, in that context, it's true, and the intuition to the contrary is merely a cognitive illusion resulting from semantic blindness. There's no indication that they also believe that, under those conditions, it's false; rather, they take the intuition that it's false to be simply misleading, albeit explicable. So, at most they believe only one of the reported intuitions.

Now, there are two objections that one might make to 'the argument from lack of belief'. First, one might claim that philosophers are simply mistaken about what they believe. Maybe, for example, they do believe the naïve comprehension axiom to be false, but hold a mistaken, higher-order belief to the effect that they don't also believe it to be true: maybe their intuition that it's true is also a belief, and they just fail to recognise this fact. Koksvik (2012) has made this point before, observing that people can clearly go wrong about what they believe, and believe contradictory things.

While Koksvik subsequently mounts his own, novel attack on strong doxasticism, I think we can save ourselves the trouble by preserving the traditional 'argument from lack of belief'. For one, while agents are indeed fallible about what they believe, unless we have specific reason to think them mistaken, or know them to be under conditions that are particularly conducive to introspective error, we should grant them more epistemic authority on this personal matter than anybody else, and give them the benefit of the doubt.

Moreover, we don't have to rely solely on contextualists', possibly mistaken, introspective awareness of their own mental states. We can also back up their belief-reports with the observation that their expressed intuitions don't always track what might be called the 'symptoms' of belief. For one, neither Cohen nor DeRose are ever prepared to actually assert the propositional contents of putatively false intuitions, but are, by contrast, willing to assert their negations. For example, of the intuitive claim that AI's conclusion is untrue, they say only that it's false, and that its negation is true. Likewise, DeRose never actually asserts the sensitivity principle, despite often calling it intuitive. Indeed, as we've seen, he explicitly declares it to be mistaken on multiple occasions. Now, generally, if S claims to not believe *p*, but reliably asserts *p*, and never its negation, we would indeed doubt the accuracy (or honesty) of their belief-report. However, if they claim to not believe *p*, and never once assert *p*, and are demonstrably willing to assert ~*p* instead, then we have every reason to trust them.

Similarly, contextualists aren't prepared to employ the propositional contents of putatively false intuitions as premises in their reasoning, only their negations. That the sensitivity principle is true, for example, is never used as the basis for DeRose's theory of knowledge, despite being an allegedly intuitive claim. In fact, DeRose will, instead, appeal to the fact that the sensitivity principle is mistaken to argue for his alternative, albeit similar, safety principle. Generally, if one's reasoning never relies on certain claims – and does, in fact, at times rely on their negations – it's a real stretch to say that those claims are included among their beliefs. In summation, we could say: contextualists aren't (always) disposed to speak and think as though their intuitions are true, so it's highly unlikely they believe (all of) them.

So, we don't have to rely only on their, potentially false introspective reportages to show that many things contextualists call intuitions are very unlikely to be counted among their beliefs.

The second objection that could be raised to 'the argument from lack of belief' against strong doxasticism is that, when contextualists continue to call *p* intuitive, even when they don't, personally, believe it, they merely mean to say that other people (the majority, other philosophers, laypersons…) believe it, or at least would believe it if asked to consider it. Thus, contextualists' intuition-talk refers to beliefs, just not always to their own.

To this, we can respond that philosophers often explicitly describe intuitions as their own. For example, in the above quotation, Bealer says: "*I* have an intuition – it still seems to *me* – that the naïve comprehension axiom of set theory is true" (my emphasis). So, he doesn't believe the axiom, but still, *himself*, has the intuition that it's true. The same can be said about the contextualists specifically. We've already seen that DeRose (2017), for example, describes the intuitiveness of AI's premises from his own perspective. He says, for instance:

> "Though *I* could always feel a strong intuitive pull toward accepting [ignorance]… *I* also felt a significant opposing intuitive push toward denying it" (p.55, my emphasis).

Cohen, meanwhile, often talks about what 'we' find intuitive (in 1988, p.116, for example). It's certainly possible that he's using this word to refer to his culture or society, and not himself. In sentences of the form 'we think it's acceptable to eat meat, but it isn't', for example, the speaker clearly doesn't intend to include themselves among the 'we'. Nonetheless, it's very likely that Cohen includes himself in this 'we', given his proclivity to employ other first-person plural pronouns, like 'us' and 'our', as well. For example, he says:

> "The non-relevance of skeptical alternatives is a datum provided by *our* intuitive judgments concerning what we know" (*ibid.*, my emphasis).

At the very least, the onus of proof is squarely on the strong doxasticist to show that, in his case, the pronouns aren't being used self-referentially.

Moreover, considering that neither Cohen nor DeRose provide any kind of support for their empirical-psychological claims about what people intuit, they very likely just extrapolate other people's intuitions from their own.

Therefore, contextualists' intuition-talk isn't just referring to the beliefs of other people.

So, the 'argument from lack of belief' against outright doxasticism holds up.

3.1.1.3 The Argument from Phenomenological Dissimilarity

Intuitions are often described as possessing a phenomenological character, that is, having an intuition 'feels' a certain way (Cappelen, 2012; Bealer, 1998). Here, for example, is an oft-cited quote from Plantinga (1993), describing the feeling of intuition as:

> "that peculiar form of phenomenology with which we all well acquainted, but which I can't describe in any way other than as the phenomenology that goes with seeing that such a proposition is true" (p.105-106).

Elsewhere, we see the phenomenology of intuition described, in similar terms, as a feeling of 'correctness' or 'rightness' (Parsons, 1995; Thompson *et al.*, 2011).

Contextualists' intuitions, specifically, are likewise phenomenal. For example, consider, again, this excerpt from DeRose (2017):

> "Though I could always *feel* a strong intuitive pull toward accepting [ignorance], I found it far from compelling, and I also *felt* a significant opposing intuitive push toward denying it" (p.55, my emphasis).

Here, the conflicted intuitions that DeRose has about 'ignorance' are explicitly described as conscious mental events that he *feels*. Likewise, when describing the intuitiveness of the sensitivity principle, he talks about a *feeling* that insensitive beliefs aren't knowledge:

> "there is another range of cases, which have their own distinctive, sometimes more tentative or sometimes conflicted '*feel*', in which we are at least pushed toward judging subjects don't know because their belief is insensitive…" (*ibid.*, p.210, my emphasis).

So, again, the intuition that the sensitivity principle is true is a conscious mental event. And, while Cohen doesn't do so quite as frequently, he also describes his intuitions as states that he 'feels' (in, for example, 1998, p.302).

Beliefs, however, are characteristically subconscious states: they usually lie 'silently' in the background, continuing to implicitly shape our behaviour without being 'felt'. I believe, for example, that grass is green, right now, even though I'm not feeling any particular way about grass whatsoever. That belief manifests in my behaviour when, for instance, I go outside and am not taken aback by the appearance of the green grass on my lawn, yet, even then, I don't feel myself holding any particular view about what hue it should be. Normally, I simply have, and act upon, that belief, without it intruding into the phenomenal character of my existence. Indeed, we might suppose that some, maybe even the majority of beliefs, never become conscious even once. Some implicit biases, for instance, might be grounded in prejudiced beliefs that never bubble to the surface of consciousness; in fact, that might be partly why such biases go undetected so easily.

This shifts the balance of evidence against strong doxasticism: the fact that beliefs are characteristically unconscious, while intuitions are straightforwardly described in phenomenal terms, suggests that beliefs and intuitions (including the intuitions cited in contextualist literature) are different kinds of state. Call this 'the argument from phenomenological dissimilarity' against strong doxasticism.

To this, one might respond that, sometimes, we do become conscious of a belief. Bear in mind that this requires more than simply knowing, or 'being aware', that the belief exists: I might, for example, learn from an implicit association test that I possess a biased belief, and still not be able to *feel* it. However, on some occasions, beliefs plausibly do, genuinely, feel a certain way. If I believe something (especially a political, moral, religious, or ideological position) passionately

enough, for example, then that belief might exert a conscious pull, or lure, upon me, urging me to action, or even be the source of a certain, mild euphoria. By contrast, if I hold a view only very tentatively, then that belief can sometimes feel 'shaky' or 'loose' in my head, or give rise to a feeling of cognitive dissonance. Perhaps it's only straightforward beliefs – those that are neither particularly passionate nor particularly hesitant (like that grass is green) – that lie silent in the background[32]. So, a token strong doxasticism could be preserved by stipulating that intuitions are just, specifically, those beliefs that are conscious.

However, the fact that intuitions are reliably described as feelings, while beliefs are (at least) almost always unconscious, still significantly weakens the case for strong doxasticism. Even if there are some conscious beliefs, it still looks, *prima facie*, as if intuitions are just a different *kind* of state from beliefs: something that, unlike belief, is essentially, not just occasionally, phenomenal. Indeed, it isn't clear what motivation we have now for thinking that intuitions are a kind of belief in the first place, given the fairly obvious difference between the phenomenology of intuition and that of (at least) almost all beliefs, which makes the aforementioned defence of strong doxasticism seem contrived.

In summation, beliefs are characteristically unconscious states. Consequently, interpreting contextualists' intuition-talk as referring to outright beliefs wouldn't properly capture the reported phenomenology of their intuitions.

3.1.1.4 The Argument from Recalcitrance

We expect people's beliefs to update when they acquire new, relevant evidence. If, for example, S believes that today is Wednesday, but then a trustworthy friend of theirs tells them it's Thursday, we would expect S to change their mind about what day of the week it is, or at least adjust their confidence levels, accordingly. Of course, people can respond badly to new evidence, such as by failing to update their beliefs in an appropriate way, or at all, and some, particularly dogged beliefs might refuse to budge even in the face of categorical evidence to the contrary. Nonetheless, on the whole, we expect people's views to change to some extent when their evidence does. Put simply, beliefs are characteristically sensitive to evidence.

This is, in part, because people can also exert a certain degree of control over what they believe. For example, if S believes *p*, and later, upon further reflection, comes to think that their evidential *corpus* contains good reason to believe ~*p*, they might conclude that their belief was mistaken

[32] I'm not entirely convinced that even these kinds of belief really have a phenomenal character; I'm simply trying to be charitable to the hypothetical advocate of strong doxasticism.

and, on that basis, adjust their view. If, for instance, S believes that today is Wednesday, but then a trustworthy friend of theirs tells them it's Thursday, S might not immediately change their mind about what day of the week it is; however, thinking about it a few minutes later, they might conclude that their friend is almost certainly right, and update their belief accordingly. At the very least, we expect it to be *possible* to bring our beliefs into alignment with our evidence through such internal deliberation. In other words, beliefs are subject to 'internal correctability': reflecting upon one's beliefs and evidence, one can normally make the former consistent with the latter, if they're sufficiently inclined.

Intuitions, by contrast, are typically thought to be insensitive to evidence. We don't expect *p* to become any more or less intuitive with changes in our evidence for or against *p*, respectively. If *p* is intuitive for us, then gaining evidence that ~*p*, while it might reduce our confidence in *p*, or even get us to stop believing *p*, if we did before, doesn't tend to make *p* any less intuitive, and *vice versa* (c.f. Huemer, 2007, p.38). The naïve comprehension axiom is a clear example of this: the axiom remains intuitive, and its denial *counter*intuitive, no matter how strong a reason one has for rejecting it. Another way to put this is to say that intuitions are 'recalcitrant' (Cappelen, 2012). A state is recalcitrant if, and only if, even if S has strong evidence against its content, it doesn't go away. This isn't to deny that there might be some, exceptional cases where intuitions do respond, to some, limited degree, to new evidence. But, while beliefs are characteristically sensitive to evidence, intuitions are characteristically insensitive to it, or, at the very least, not nearly as sensitive to it as beliefs are.

This is partly because, by contrast to our beliefs, people also can't exert control over their intuitions. If S intuits *p*, then they cannot stop themselves from doing so, through any degree of introspective reflection or effort, even if they take themselves to have overwhelming evidence that *p* is false. That is, S can't bring their intuitions into agreement with their evidence, even if they're so inclined, and thus intuitions don't display 'internal correctability'. Again, we might recall how Bealer and other philosophers continue to feel the intuition that the naïve comprehension axiom is true, for example, even when they take themselves to have decisive reason against it; presumably, if it were possible for them to shed that false intuition, or otherwise bring it into line with their evidential *corpus*, they would have done so, if only to relieve themselves of the cognitive dissonance.

So, beliefs are typically sensitive to evidence, in part because they can be adjusted through internal reflection, while intuitions are insensitive to evidence, partly because people don't have any equivalent control over what they intuit. Intuitions and beliefs are, therefore, quite different kinds of thing. Call this 'the argument from recalcitrance' against strong doxasticism.

This argument clearly applies to contextualists' intuitions, specifically. There are several things that contextualists report finding intuitive, but take themselves to have decisive evidence against: the intuition that S doesn't know they're a BIV, even by ordinary epistemic standards, the intuition that 'scepticism' is false when one is presented with AI, the intuition that sensitivity is a general condition on knowledge, and so on. If these intuitions were beliefs, we would expect them to disappear once the contextualists had good arguments against them. Because the contextualists report still having them, they're insensitive to evidence in a way that beliefs usually aren't. Furthermore, contextualists don't seem able to bring their intuitions into line with their evidence through internal reflection or effort; again, if such a thing were possible, we would expect them to have done it before now.

So, to put it simply, the fact that contextualists still have intuitions that they take to be false is yet further reason not to interpret contextualists' intuition-talk as simply referring to outright beliefs.

3.1.1.5 The Argument from Explanatory Redundancy

Sometimes, when S intuits *p*, they'll subsequently come to outright believe *p*, on the basis of that intuition (Bengson, 2015, p.727; Conee, 1998, p.850; Huemer, 2007, p.31). In such cases, we can say that S' intuition grounds, or explains, their corresponding belief. That explanation seems informative: it tells us something about how and why S' belief came about. However, if intuitions were just outright beliefs, the same explanation would amount to saying no more than: S believes *p* because S believes *p*. Such an account sounds, by contrast, uninformative; we don't learn anything new about S' belief, or how and why they formed it. Thus, strong doxasticism renders intuitions explanatorily redundant, where they shouldn't be. Call this 'the argument from explanatory redundancy' against strong doxasticism.

Again, this argument also applies to contextualists' intuitions, specifically. There are, as we've seen, several instances of contextualists defending some proposition, *p*, on the grounds of its intuitiveness. In such cases, we can assume that they believe *p* – since they're trying to convince their readers that it's true – and that their intuition that *p* at least partly grounds, or explains, that belief. For example, it's highly plausible that Cohen accepts certain claims about what S does and doesn't know in certain contexts because he intuits S to know or not know things in said contexts, and that he presents these intuitions to others as a means of conveying his reasons for belief, in the hopes that they will accept *p* on the same grounds. Again, these facts sound informative: it tells us something about why the contextualists believe what they do, that is, because they have certain intuitions about things. But, if their intuition that *p* was just their belief that *p*, this would make for a vacuous description. Thus, their intuitions don't appear to be beliefs.

We might also observe that, for similar reasons, strong doxasticism entails a deeply uncharitable reading of the contextualists' argumentative method. They do, as we've seen, argue for *p* on the basis of its intuitiveness; if this merely amounts to arguing for *p* on the basis that they believe *p*, that's an extremely weak argument. So, a charitable reconstruction of their terminology shouldn't equate their intuitions with beliefs.

To summarise, on the basis of these four arguments, we shouldn't interpret contextualists' intuition-talk as referring to outright beliefs. To close this section, I'll consider how we might reconcile this with the aforementioned textual evidence of strong doxasticism.

3.1.1.6 Accommodating Problematic Language

To reiterate, despite some superficial evidence in its favour, the interpretation of contextualists' intuition-talk as referring to outright beliefs is implausible. What, then, should we say about that aforementioned textual evidence to the contrary – about their belief-talk, judgement-talk, and attribution-talk?

Well, as previously mentioned, using belief-talk interchangeably with intuition-talk doesn't occur in any of Cohen's work, and only happens once in DeRose – the excerpt I discussed earlier is, in fact, the only clear instance where he uses belief-talk *in place of* intuition-talk. We could, therefore, say that, on this one occasion, DeRose just slipped into using belief-talk, although misleadingly. Because he's never thought to provide a considered definition of intuitions, it wouldn't be at all surprising if his descriptions of them were somewhat unreliable at times.

Judgement-talk, however, is harder to dismiss. Again, this appears only once (at least in a manner where it clearly seems to be describing intuitions) in Cohen. However, it does, as we've seen, manifest quite prolifically in DeRose. If DeRose's intuitions aren't beliefs, then why does he so often use judgement-talk and intuition-talk interchangeably? Well, one, charitable possibility is that, in DeRose, judgement-talk has simply taken on a non-doxastic meaning. After all, while using judgement-talk to describe beliefs – and, in particular, the formation of beliefs – is certainly the norm among contemporary epistemologists, it's not used in quite such a manner everywhere, not even in philosophy. In moral philosophy, for instance, 'moral judgements' doesn't necessarily mean moral beliefs, or the formations of moral beliefs. Rather, it can also refer to speech acts of the form 'murder is wrong'. Indeed, some – the non-cognitivists – don't even hold that such 'judgements' *express* beliefs. Van Roojen (2016), for example, says:

> "many non-cognitivists hold that moral judgments' *primary* function is not to express beliefs" (§1.1, emphasis in original).

Perhaps DeRose is among a crowd who have similarly taken to using judgement-talk in a more liberal, non-doxastic manner. Maybe, for him, 'intuitive judgements' just means 'whatever kind of things intuitions are', which isn't necessarily beliefs or the acts of forming them.

That said, ultimately, it's plausible that DeRose's presentation of the argument for EC is just confused. Again, the above suggestion is meant to be charitable.

Finally, regarding Cohen's aforementioned attribution-talk, I'll say more about this particular mode of speech later in this chapter (§3.3), wherein I discuss the possibility that contextualists' intuitions are facts about ordinary language. There, I'll argue that the intuitions which Cohen is endeavouring to account for with his contextualist theory of knowledge are intuitions *about* certain knowledge 'attributions', or their content, rather than the attributions themselves. If we interpret his attribution-talk as referring to outright beliefs, the relevant intuitions are, therefore, *about* outright beliefs, or their content, not the beliefs themselves. But again, I'll discuss this in more detail later.

The fact that belief-talk, judgement-talk, and attribution-talk appear so pervasively in contextualist literature, as apparent synonyms for intuition-talk, reminds us of the need to reconstruct their argument in clearer terms: as it stands, their description of the alleged 'intuitive data' is incoherent and confusing. Nonetheless, their proclivity to use this terminology does suggest that intuitions at least somewhat resemble beliefs, and are perhaps even connected to them in an important way, and my novel account of intuitions, which I develop in Chapters 4-5, will do justice to that relationship.

To recapitulate, I've argued that we shouldn't reconstruct the contextualists' intuition-talk as referring to outright beliefs. In the next section, I'll argue that we also shouldn't equate their intuitions with partial beliefs (or credences).

3.1.2 Against Moderate Doxasticism

Instead of identifying intuitions with outright beliefs, an alternative doxastic approach is to identify them with partial beliefs, a possibility also discussed in Koksvik (2012). Call this view 'moderate doxasticism'. To elaborate, S' partial beliefs, or 'credences', are how confident they are that propositions are true. S might, for example, think that it's 60% (or 0.6) likely that it'll rain tomorrow, and 40% (or 0.4) likely that it'll be sunny. These are, therefore, their confidences, or partial beliefs, in those propositions. Accordingly, S' credences correlate (at least roughly) with, and manifest in, their betting habits: S' credence in *p* determines how much money they'd be

willing to risk on *p* being true, since it's how much of a gamble they take it to be (Jeffrey, 1965, p.60, c.f. Christensen, 2004, Chap. 5; Ramsey, 1931; Skyrms, 1980).

Now, a type moderate doxasticism, that conflated all credences with intuitions, and *vice versa*, would, like type strong doxasticism, massively overgenerate; clearly, there are very many propositions that we assign at least some credence to, including, presumably, all those that we deem likely enough to outright believe, and not all such propositions are intuitive – again, many that we outright believe are not, for example. The likes of Bealer, for instance, feel very confident that the negation of the naïve comprehension axiom is true, but that negation remains paradigmatically *counter*intuitive. Thus, the only plausible moderate doxasticism is, likewise, of the token, not type variety, to wit, a view to the effect that all intuitions are token credences, but not all token credences are intuitions.

In this section, I'll first show that token moderate doxasticism, like token strong doxasticism before it, has a modicum of textual evidence in its favour, and is therefore worth taking seriously. Then, I'll present two arguments against interpreting contextualists' intuition-talk, specifically, as referring to partial beliefs.

3.1.2.1 Textual Evidence of Moderate Doxasticism

There's some textual evidence in favour of interpreting contextualists' intuition-talk as referring to partial beliefs, or credences. For example, in both Cohen and DeRose's work, intuition-talk often appears alongside, and is used interchangeably with, plausibility-talk. Looking first at DeRose (2017), we see here how he vacillates between calling the sensitivity principle 'intuitive' and calling it 'plausible':

> "[C]onsider why [sensitivity] is an attractive explanation in the first place… Much of that attraction… comes from the initial *plausibility* of [sensitivity's] generalization: Pointing out that someone would have held a certain belief even if it had been false *intuitively* seems a very good reason for denying that they know what they believe" (p.182, my emphasis).

Likewise, when describing what he elsewhere calls our 'intuitions' toward AI, DeRose sometimes instead says things like "each of its premises is initially *plausible*" (*Ibid.*, p.2, my emphasis) and criticises non-contextualist responses to the argument for having to give up at least one of these intuitive claims by saying that, if we follow them, "something *plausible* has to go" (*Ibid.*, my emphasis). Moreover, when discussing the intuitiveness of 'ignorance' and 'closure', he says:

> "We must explain how two premises that together yield a conclusion we find so incredible can themselves seem so *plausible* to us" (*ibid.*, my emphasis).

Meanwhile, we find similar patterns of speech in Cohen, who likewise fluctuates between intuition-talk and plausibility-talk. Recall, for example, that his contextualism promises to explain, among other things, "the *intuitive* appeal of the deductive closure principle for knowledge" (2005, p.205, my emphasis). However, he also, on occasion, describes the appeal of CPK in terms of its 'plausibility':

> "The paradox arises given *a plausible closure principle* that entails that Smith knows he will never get rich only if Smith knows he will lose the lottery" (*Ibid.*, p.200, my emphasis).

In fact, when describing the Harman-Vogel paradox, Cohen casually vacillates between describing its constitutive propositions as 'intuitive' and calling them 'plausible', such as here:

> "We should view [closure], [ignorance], and [~scepticism] as constituting a paradox-a set of inconsistent propositions all of which have considerable independent *plausibility*… But a satisfying resolution requires an explanation of why the paradox arises-an explanation of why we have the *intuitions* that saddle us with the paradox" (1988, p.94, my emphasis).

Additionally, there are instances where DeRose, at least, chooses to describe the intuitive data to be explained by contextualism in terms of how 'persuasive' the premises of AI are:

> "[W]e seek to *explain the persuasiveness* of AI, at least in large part, by claiming that the presentation of the skeptic's argument has at least some tendency to put into play the very standards at which we don't count as knowing that we have hands" (2017, p.64).

Finally, Cohen will sometimes describe the intuitive data in terms of our 'confidence'. For example, when using his contextualism to explain the counterintuitiveness of 'scepticism', he sometimes says that his view:

> "explains our *confidence* in the truth of our everyday attributions of knowledge" (1988, p.96, my emphasis).

So, contextualists sometimes describe the intuition that *p*, at other times, as finding *p* 'plausible' or 'persuasive'. This suggests that, for them, intuiting *p* just means partially believing *p*: assigning sufficient likelihood to *p* to at least find it plausible or persuasive, to some degree. Thus, contextualists' intuition-talk might just be referring to partial beliefs. Indeed, we even see, in the

last quote, how Cohen sometimes even describes one of the relevant intuitions as our 'confidence' in *p*, which arguably refers, more straightforwardly, to our credence.

So, there's a modicum of textual evidence to suggest that contextualists' intuition-talk refers to partial beliefs.

Token moderate doxasticism does circumvent some of the problems associated with token strong doxasticism. For instance, token moderate doxasticism, unlike its counterpart, is compatible with the observation that not all intuitions are believed outright by their subjects. To elaborate, S can hold a partial belief in *p* without outright believing *p*, since, for example, they might not take *p* to be likely *enough* to commit themselves to the truth of that proposition. They might not, for instance, be willing to straightforwardly assert *p*, or use *p* (*simpliciter*, rather than the probability of *p*) as a premise in their practical reasoning. Thus, the interpretation of contextualists' intuition-talk as referring to credences is consistent with them failing to outright believe, for example, that the sensitivity principle is true: perhaps DeRose 'intuits' this in the sense that he finds it somewhat plausible, but doesn't find it plausible enough to outright believe it.

However, there are still several reasons to reject this particular reconstruction of contextualists' intuition-talk. For one, token moderate doxasticism inherits some of the same problems as token strong doxasticism. For example, we expect our partial beliefs, like our outright beliefs, to change with the addition or subtraction of relevant evidence. If, for instance, S is 60% confident that today is Wednesday, and a trustworthy friend tells them that it's Thursday, then we expect S to at least become less confident that it's Wednesday. And, as with outright beliefs, this is partly because S can exercise a certain degree of control over their confidence levels; reflecting on their reasons for belief can lead S to internally adjust their credences, and bring them into agreement with their evidential *corpus*, if they're so inclined. Again, this isn't to deny that we can have recalcitrant credences – we can, just as we can have recalcitrant beliefs – but, characteristically, partial beliefs are sensitive to evidence and internally correctible. So, they likewise differ from contextualists' intuitions in this respect. Partial beliefs also differ from contextualists' intuitions phenomenologically, since, like outright beliefs, our confidences are usually unconscious; it doesn't normally feel like anything to think that, for example, it's 60% likely to rain tomorrow. Credences like this, just like beliefs, sit 'silently' in the background, implicitly shaping, and manifesting in, our behaviour, such as our betting habits, without being consciously experienced, or 'felt', by their subjects.

Below, I also offer two, additional arguments against moderate doxasticism specifically.

3.1.2.2 The Argument from Disproportionality

Partial beliefs come in degrees: S can be only slightly confident that *p* is true, very confident, or anything in between. Intuitions are likewise gradable: the intuition that *p* can vary from very weak to very strong. DeRose (2017, p.54), for example, observes that our intuitions about 'ignorance' are much weaker and more conflicted than our comparatively strong and stable intuitions about 'closure'. If intuitions were just partial beliefs, then we would expect stronger intuitions to straightforwardly correspond with higher credences; if 'intuiting' was just a matter of partially believing, then a *stronger* intuition should just be a *higher* confidence level (and *vice versa*). However, on the contrary, the strength of one's intuitions doesn't reliably coincide with the strength of one's credences. For example, there are cases where philosophers outright believe against even strongly felt intuitions. Some, for instance, deny CPK, despite finding it *extremely* intuitive (see, for example, Dretske, 2005). In such cases, we can assume that they still hold at least a partial belief that CPK is true, if they're at all intellectually modest and recognise some chance of their own view being false. However, that partial belief would, at best, be quite weak, and certainly not proportionate to the strength of the reported intuition. Indeed, if their confidence level in CPK was equivalent to its intuitive strength, they'd likely just outright believe the principle. Moreover, consider that, when people like Bealer reject the naïve comprehension axiom, for instance, they likely assign it no credence whatsoever, since, on the grounds of Russel's paradox, they take it to entail a contradiction. In other words, it's possible to intuit *p*, to some degree, without partially believing *p*, to any degree (c.f. Earlenbaugh and Molyneux, 2009). In brief, we can say that intuitions don't track partial beliefs – the strength of the two phenomena isn't proportional. Thus, intuitions and partial beliefs aren't the same thing. Call this 'the argument from disproportionality' against token moderate doxasticism.

This argument also works against the interpretive thesis that contextualists' intuitions, specifically, are just partial beliefs, since they also hold outright beliefs that contradict their strong intuitions. For instance, recall that DeRose and Cohen both firmly believe that S does know, by ordinary epistemic standards, that they're not a BIV. Nonetheless, they both find it 'strongly' or 'significantly' intuitive that S doesn't know this, by any standards (see, for example, DeRose, 2017, p.55). Again, the contextualists presumably have some, rational doubt about their own theories of knowledge, and thus assign some probability to the contrary, but surely that doubt isn't proportional to the 'strong' intuition that they report experiencing to the effect that not being a BIV is unknowable. Indeed, if their confidence in the claim that S can't know they're not a BIV *was* equivalent to the strength of their reported intuition, they'd likely just accept that S can't know this.

In summary, contextualists' intuitions aren't always the same strength as their credences, so their intuitions aren't just partial beliefs.

3.1.2.3 The Argument from Irrationality

It isn't irrational for S to have a strong intuition that *p*, even if they outright believe ~*p* (c.f. Koksvik, 2012, p.56-7). For instance, recall, once more, that many philosophers, like Bealer, outright believe that the naïve comprehension axiom is false, because they have great evidence against it (they're familiar with Russell's paradox, and know it to entail a contradiction). Nonetheless, it doesn't seem irrational for those same philosophers to find the axiom very intuitive (which, as we've seen, they do). Indeed, it would seem quite excessive to criticise them simply for feeling the intuition. In this respect, intuitions are, again, somewhat akin to perceptual experiences; it likewise isn't irrational for S to have a perceptual experience that contradicts their outright beliefs, even when that experience is especially vivid. For instance, it isn't irrational for S to continuing seeing a submerged stick as bent when they outright believe, and have every reason to believe, that it's straight. By contrast, holding a strong partial belief that *p*, while outright believing ~*p*, *does* seem irrational. If, for instance, I'm 80% confident that it will rain tomorrow, but outright believe that it won't, then I'm rationally criticisable. At the very least, what propositions I believe should line up with how likely I take them to be.

This difference might be, in part, due to the aforementioned recalcitrance[33] of intuition and perceptual experience. S can't, for instance, stop seeing a submerged stick as bent, no matter how sincerely they try. Plausibly, this means that having that perceptual experience isn't irrational, even if they believe, and have every reason to believe, that it's misleading. If they can't not have the experience, then it can't be the case that they *ought* not to have it, and so having it isn't violating any robust epistemic norms, whatever beliefs or evidence they happen to possess. Again, intuitions are likewise recalcitrant, as previously discussed. Thus, as with perceptual experiences, having intuitions that contradict your outright beliefs isn't criticisable. S can't not have the intuitions that they do, so it can't be irrational for them to have them, whatever other evidence or beliefs they have. By contrast, to reiterate, S does have some agency over their partial beliefs. They can usually bring them into harmony with their evidential *corpus*, with one another, and with their outright beliefs, if they recognise any inconsistencies, and are so inclined –

[33] While being recalcitrant and not displaying internal correctability are, as I've defined them earlier, different properties, a state that isn't internally correctible will most likely be recalcitrant, and a state that is recalcitrant will probably be so, at least party, because it isn't internally correctible. The two properties are closely tied in such a way that I will sometimes treat them interchangeably, though never where this would become a problem.

credences are 'internally correctable' phenomena. Thus, S can be rationally criticised if they have strong credences that contradict their outright beliefs. It's possible for them to make their credences coherent with their outright beliefs, so we can criticise them for failing to do so.

Whatever we take the reason to be, however, having intuitions that contradict your outright beliefs isn't rationally criticisable, while having strong credences that contradict your outright beliefs is. Therefore, intuitions and partial beliefs are subject to different epistemic norms, suggesting, once again, that they're different kinds of state.

Moreover, consequently, if contextualists' intuition-talk was referring to credences, this would make them irrational, for the reasons given. After all, they do, as previously stated, have strong intuitions that contradict their outright beliefs. For example, they have a strong intuition that S can't know they're not a BIV, by ordinary epistemic standards, and outright believe that they can. If this intuition were just a strong credence, they'd be rationally criticisable for not aligning their outright beliefs with their confidences. Now, contextualists, like all philosophers, are, of course, more than capable of being irrational, but the principle of interpretive charity cautions us against any interpretation of contextualist literature that paints them as excessively so. Thus, we also have principled reasons to reject moderate doxasticism about contextualists' intuitions.

So, the interpretation of contextualists' intuition-talk as referring to partial beliefs, or credences, shares problems with the outright belief reading, and has flaws of its own.

In conclusion, we shouldn't interpret contextualists' intuition-talk as referring to outright beliefs, nor to partial beliefs. In the next section, I'll argue against another minimalist account of intuitions, known as 'dispositionalism'.

3.2 Against Dispositionalism

Another minimalist account of intuition says that intuitions are neither outright beliefs, nor even partial beliefs, but rather dispositions to believe, or 'doxastic dispositions'. Call this theory 'dispositionalism about intuitions', or just 'dispositionalism'[34]. Such a view has been advocated by, for example, Sosa (2006), Lynch (2006), and Earlenbaugh and Molyneux (2009). In this section, I'll argue that, although contextualists' intuitions likely ground doxastic dispositions, they aren't

[34] We could also call this 'weak doxasticism', since it still ties intuitions to beliefs, just less directly than either of the doxastic positions discussed previously. However, since it's not a view according to which intuitions *are* beliefs (whether outright or partial), I've chosen not to do this.

identical with them[35]. Thus, I'll reject a dispositional account of contextualists' intuitions, specifically.

Before proceeding, however, we need to carefully distinguish between what I will herein call 'being disposed to Φ' and 'feeling tempted to Φ'. Being disposed to Φ just means reliably Φ-ing. This state is often expressed conditionally, in the form 'S has the disposition to Φ if *C*', meaning, S would (reliably enough) Φ if *C*. I, for example, have the disposition to take an umbrella out with me *if it's raining*, meaning, I would (reliably enough) take one if it were. *C* denotes the 'manifestation conditions' of the disposition to Φ: the conditions under which the disposition will manifest, or, under which S would Φ. Note, however, that S need not invariably Φ if *C* to have the disposition to Φ if *C*. Indeed, we can think of the frequency with which S does Φ if *C* as the 'degree' to which they're so disposed. We can clarify this further utilising a very basic 'possible world' semantics. There are logically possible worlds where *C* obtains, *C*-worlds, and, if S Φ-ies in enough of those *C*-worlds, then S has the disposition to Φ if *C*. The number of *C*-worlds in which S Φ-ies is the degree to which they are so disposed.

Feeling tempted to Φ, on the other hand, means experiencing a conscious temptation to Φ. This state might also be expressed conditionally, in the form 'S feels tempted to Φ if *C*'. I, for example, feel tempted to eat chips *if I smell them*. Again, the frequency with which S experiences this wanting if *C* is the 'degree' to which they feel tempted. Someone who very often experiences the temptation to eat chips when they smell them is, in a sense, more tempted to eat chips under those conditions than someone who only rarely feels this. However, temptations also vary in 'intensity': a subject can *feel* the temptation to Φ more or less strongly and/or vividly. I, for example, reliably feel *quite strongly* tempted to eat chips when I smell them, but someone else, who finds them less delicious, might feel the temptation much less intensely, even if just as reliably. *Ceteris paribus*, the stronger and/or more vivid the temptation to Φ, the harder it'll be for S to resist actually Φ-ing.

Now, there's an obvious connection between being disposed and feeling tempted, in that the latter will sometimes ground (cause, or explain) the former. If S is disposed to feel tempted to Φ if *C*, and that temptation is intense enough to typically overcome any resistance, then S is also disposed to Φ if *C*, for that very reason (since, if *C*, then S would surrender to temptation and Φ). If, for example, I'm reliably tempted to eat chips when I smell them, and that temptation is strong enough to typically overcome any resistance, then I'm also disposed to eat chips when I smell

[35] For simplicity, I'll understand doxastic dispositions as, specifically, dispositions to form outright beliefs. However, my criticisms could easily be modified to apply to a version of dispositionalism that equates contextualists' intuitions with dispositions to form partial beliefs instead.

them, for that very reason. Nonetheless, dispositions and temptations are distinct phenomena. For one, they can exist independently. For example, I might be disposed to take an umbrella out with me when it's raining, but not (and not because I) feel tempted to do so. I might just do this unconsciously, as a force of habit. Here, the thing I'm disposed to do isn't mediated, motivated, or accompanied by a conscious temptation to do it. Conversely, a lifelong vegetarian, for example, might reliably feel tempted to eat meat on certain occasions, but manage to systematically resist that urge, and never once actually do so, such that they're not *disposed* to eat it. Here, the temptation to Φ doesn't ground, and isn't accompanied by, a disposition to Φ. Moreover, while being disposed to Φ is a 'standing' state – one that manifests if *C*, when S actually Φ-ies, but which is present in S even when they're not Φ-ing , since it's still true of them that they *would* Φ if *C* – feeling disposed to Φ is an occurrent, conscious mental event: when *C* doesn't obtain, and S isn't occurrently, consciously feeling tempted to Φ, the temptation is simply absent. In other words, *C* isn't merely the conditions under which a temptation will manifest, but the conditions of its very existence.

Yet, in the extant literature on dispositionalism, being disposed to believe hasn't been properly distinguished from feeling tempted to believe. At times, dispositionalists seem to be advocating the view that intuitions are dispositions in the sense that I describe 'being disposed' above, while, at others, their position seems closer to saying that intuitions are more like what I call 'feeling tempted'. Sosa (2007c), for example, says:

> "to intuit that p is to be attracted to assent simply through entertaining that representational content. The intuition is rational if and only if it derives from a competence" (p.101).

Here, an intuition is an 'attraction' to assent to a proposition. However, this attraction could be read as either a disposition or a temptation. On the one hand, Sosa ties 'rational' intuitions to conceptual 'competence' (a fairly common view that I'll also discuss later, c.f. Ludwig, 2010, p.435, 2007, p.135; Bealer, 1996, 2000; Goldman, 2007; Richard, 2014; Kornblith, 1998) and, elsewhere, defines competences as dispositions:

> "A competence is a certain sort of disposition to succeed when you try… it is a competence in part because it is a disposition to succeed reliably enough when one makes such attempts... It is thus tied to a conditional of the form: if one tried to Φ, one would (likely enough) succeed" (Sosa, 2015, p.96).

So, conceptual competences are dispositions to apply concepts 'successfully', or, we could say, 'correctly': S is competent with concept X if, and only if, S would, often enough, accurately apply X

if they tried. Rational intuitions, Sosa says, 'derive' from such competences, and might, therefore, be considered further dispositions that are grounded in them. If S is disposed to reliably apply concept X, and this alone grounds, or explains, S' disposition to believe that *f* is X, then the latter disposition is a 'rational intuition'. If, for example, I'm disposed to reliably apply the concept 'green', and this conceptual competence alone grounds my disposition to believe that grass is green, then the latter disposition is a 'rational intuition'. We can suppose that irrational intuitions are likewise dispositions to believe that *f* is X, but aren't, by contrast, tied to a more general disposition to apply the relevant concept, X, correctly. So, for Sosa, intuitions, or 'attractions to assent' could, therefore, be dispositions, in the sense that I defined them earlier: all intuitions are dispositions to believe things of the form '*f* is X', and are rational if, and only if, they're grounded in a more general competence with the relevant concept, X. Here, when S has an intuition that *f* is X, they don't 'feel tempted' to believe that *f* is X. Rather, it's just the case that, if they reflected upon *f* and X, they would, reliably enough, (rationally/competently or irrationally/incompetently) form that belief.

On the other hand, though, Sosa (2009, section 1) also says that attractions to assent, or 'intuitions', are akin to 'desires'. Desiring to do something, however, sounds more like feeling tempted to do it than being disposed to do it. Indeed, it's consistent with desiring to Φ that one *isn't* disposed to Φ. One might, for instance, be unable to (ever) do the thing they desire, Φ, or (always) have a stronger, overriding desire to ~Φ. Thus, if intuiting *p* means desiring to believe *p*, then it's closer to a temptation than a disposition. Moreover, Sosa sometimes describes intuitions as, specifically, *conscious* inclinations to assent (for example, on 2007a, p.47). This makes them sound like conscious, occurrent mental events, like temptations, rather than unconscious, standing states, like dispositions. He likewise calls them 'pulls' (*Ibid.*), which again suggests that they're something like temptations to believe: being consciously drawn or lured in the direction of belief. So, Sosa's dispositionalism is ambiguous between a 'being disposed' and 'feeling tempted' account of what intuitions are.

In this section, I'll only criticise the view that contextualists' intuitions are *being disposed* to believe, which, for the sake of what immediately follows, I'll just call 'dispositionalism'. This is because that's the version of dispositionalism that is inarguably minimalist, since it reduces intuitions to a sub-class of another state, namely dispositions *simpliciter*. Feeling tempted to believe, on the other hand, is arguably a *sui generis* phenomenon, as some philosophers, including Sosa, have suggested. Accordingly, I discuss the possibility that contextualists' intuitions are doxastic temptations in a later section of the thesis (§4.4), after I've turned my attention away from extant minimalist accounts. Ultimately, however, I conclude that contextualists' intuitions are neither dispositions to believe nor temptations to believe.

A further thing to note is that S has a corresponding disposition to believe any proposition that they do, in fact, believe, or have ever believed. This is because, if S believes, or has believed, *p*, then there are, demonstrably, some conditions, *C*, under which S (at least sometimes) believes *p* and, thus, S is disposed (to some degree) to believe *p* if *C*. In other words, we just have to adequately specify the manifestation conditions of the corresponding disposition, to capture the circumstances under which the proposition is, or has been, believed. Therefore, under a type dispositionalism, according to which all doxastic intuitions are intuitions, and *vice versa*, any believed proposition whatsoever would have a corresponding intuition. This is clearly false; again, it's possible to believe unintuitive things, such as the negation of the naïve comprehension axiom. This applies equally well to the contextualists, specifically. For example, the claim that knowledge-talk is semantically context-sensitive is, recall, meant to be counterintuitive, because of our 'semantic blindness' (see, for instance, Cohen, 2005, p.208). Nonetheless, this is something that contextualists, at least, believe. Accordingly, the only plausible dispositionalism is, once more, of the token, not type variety, to wit, a view to the effect that all intuitions are token doxastic dispositions, but not all doxastic dispositions are token intuitions.

With these preliminaries out of the way, I'll now discuss the possibility that contextualists' intuition-talk refers to (token) dispositions to believe. I'll start by presenting the textual evidence in favour of such an interpretation, which warrants us taking the hypothesis seriously. Then, I'll present a number of objections to this view.

3.2.1 Textual Evidence of Dispositionalism

There are some textual indications that, for the contextualists, intuitions are just dispositions to believe. Here, for example, the intuition that we know ordinary things (such as that we have hands) (or 'Os'), when they're considered in isolation, is contrasted against the intuition that we don't know the negations of sceptical hypotheses (such as that we're not a BIV) (or 'not-Hs'). The latter intuition, however, is now described as a strong 'tendency' to say or think that we don't know:

> "[W]hen we consider our new Os by themselves, they seem like things we know to be the case, but when we consider the not-Hs, there is at least *a strong tendency to say/think* that we don't know them to be the case" (DeRose, 2017, p.154, my emphasis).

To my ears, having a tendency to Φ sounds more like being disposed to Φ than merely feeling tempted to Φ: it sounds like, reliably enough, Φ-ing. So, the intuition that we don't know not-Hs could just be a disposition to say or think (or believe) that we don't.

Elsewhere, DeRose says that that the safety principle, like the sensitivity principle, is supported by its intuitiveness. However, he describes the support for the former principle, specifically, as the fact that we 'often judge' subjects to not know *p* when their belief that *p* could have easily been false (or is 'unsafe'):

> "Often we judge subjects not to know something because, though their belief is true, it very easily could have been false…" (*ibid.*, p.209-210).

So, for DeRose, the intuition supporting the safety principle could just be a disposition to judge, or believe.

Also, while Cohen sometimes describes the Harman-Vogel paradox as an inconsistent set of 'intuitions', he elsewhere describes it as an inconsistent set of 'inclinations to accept', as he does here:

> "we are faced with a paradox. We are *inclined* to accept each member of a set of propositions we know to be inconsistent. What we seek is a way out of the paradox-a resolution of our inconsistent *inclinations*" (1988, p.113, my emphasis).

So, perhaps the intuitions constituting the Harman-Vogel paradox are just inconsistent inclinations to accept, or believe, propositions. Cohen similarly describes the intuition that 'scepticism' is true (not only for us, in the context where we're reflecting on AI, but elsewhere as well) as an 'inclination to deny' everyday knowledge claims:

> "And when we are in this frame of mind, not only do we have a strong *intuitive inclination* to deny that we know, we have the same *inclination* to deny that others know" (2005, p.207, my emphasis).

To my ears, having an 'inclination' to Φ again sounds more like being disposed to Φ than merely feeling tempted to Φ: it sounds like, often enough, Φ-ing. So, for Cohen, our intuitions could just be dispositions to 'accept' or 'deny' propositions, or, to believe or disbelieve them.

In the next sub-section, I'll argue that this textual evidence is misleading, and contextualists' intuition-talk doesn't refer to doxastic dispositions.

3.2.2 The Argument from Lack of Disposition

Since S being disposed to believe *p* if *C* is simply a matter of S believing *p* if *C* – where the frequency of said belief under those conditions is the strength of their disposition – whether S is disposed to believe *p* if *C*, *at all*, is determined simply by whether S does, on some occasions,

believe *p* if *C*. If S doesn't ever believe *p* if *C*, then S isn't at all disposed to believe *p* if *C*. And, if S doesn't ever believe *p* under *any* conditions, then S isn't disposed to believe *p simpliciter*. In other words, just as the occurrence of a belief in *p* demonstrates the existence of a disposition to believe *p*, the absence of a belief in *p* demonstrates the absence of any such corresponding disposition. Accordingly, because, as we've seen, contextualists can intuit propositions that they don't believe, they can have the intuition that *p* without having a corresponding disposition to believe *p*. Therefore, contextualists' intuitions aren't identical with any token doxastic dispositions. Call this 'the argument from lack of disposition' against a token dispositionalist reading of contextualists' intuition-talk.

Now, one could object to this argument that S failing to Φ on some, particular occasions doesn't prove, *ipso facto*, that S isn't disposed to Φ, to some degree, under some conditions. For one, S might be disposed to Φ if *C* to some degree, but still not always Φ if *C*. I might, for example, be disposed to take an umbrella out with me if it's raining, but still not always do so. I might sometimes forget. Alternatively, S might be disposed to Φ, not if *C*, but under different circumstances, *C**. I might, for example, be disposed to take an umbrella out with me when it's raining hard, but not when it's raining lightly. Either way, the fact that S doesn't Φ, in a particular time and place, is insufficient to prove that S isn't disposed to Φ, to some degree, under some conditions. Thus, the mere fact that contextualists sometimes fail to believe the propositions they intuit – when, for instance, they're thinking about contextualist solutions to the Harman-Vogel paradox – doesn't prove that they're not disposed to believe them, to some degree, under some conditions. So, their reported intuition could still, conceivably, be such a doxastic disposition.

Moreover, the disposition to Φ if *C*, even if it otherwise *would* manifest on a particular instance of *C*, can be 'blocked', and thus prevented from doing so. As Koksvik (2012) puts it:

> "[W]here *x* is in C but does not Φ *does not show* that *x* is not disposed to Φ in C. For *x*'s having the disposition to Φ in C is consistent with something *blocking* that disposition from manifesting on a particular occasion" (p.82, emphasis in original).

For example, even if I'm disposed to take an umbrella out with me when it's raining, and would otherwise have done so on this, particular rainy day, if I have misplaced my umbrella, and am unable to locate it, this will prevent me from taking it out, and thus 'block' that disposition from fully manifesting. The difference here between S' disposition to Φ if *C* not manifesting even when *C* simply because S is only disposed to *sometimes* Φ if *C*, and the disposition to Φ if *C* being blocked, as I'm using these terms, is that, in the former case, the disposition simply isn't 'triggered' or 'initiated', while, in the latter case, it is: S attempts to Φ, and is prevented, so the disposition to Φ fails to completely manifest in S actually Φ-ing. Indeed, S being disposed to Φ is,

in principle, compatible with that disposition being systematically blocked on all the occasions where it would otherwise manifest, and S failing to Φ at all. So, even if the contextualists *never* believe an intuited proposition, this doesn't prove that they aren't disposed to believe it, and their intuition could still, in principle, be identical with such a, systematically blocked, doxastic disposition.

My reply to this is that the best and most obvious evidence of S having a disposition to Φ (*simpliciter*, or, specifically, if *C*) is still S actually Φ-ing (at all, or, specifically, when *C*). That is what we'd normally expect to see if S really had such a disposition. So, the less reason we have for thinking that S actually Φ-ies (at all, or, specifically, when *C*), the less reason we have to believe that S is disposed to Φ (at all, or, specifically, when *C*). Indeed, when there's no evidence whatsoever of S ever actually Φ-ing, it's much more rational to simply *dis*believe that S has any disposition to Φ. As they say, absence of evidence, after you've looked, is evidence of absence. If, for example, you observe my behaviour very closely over an extended period of time, and never once do you see me take an umbrella out with me, even when it's raining, then it's reasonable for you to surmise that I don't have a disposition to take an umbrella out with me when it's raining. Now, the contextualists have at least some intuitions that they've never actually believed. For example, closure-based sceptical arguments, like AI, are meant to give us the intuition that we don't know whether we have hands, even by ordinary epistemic standards (recall Cohen, for instance, describing the sceptical intuition that our everyday knowledge claims, like this one, are mistaken). However, it's most unlikely that the likes of Cohen and DeRose have ever believed that they don't know (by ordinary standards) whether they have hands. In fact, Cohen (1988) explicitly states that, despite AI making such scepticism seem intuitive, "what we are confronted with is *not* an argument that forces us to be skeptics" (p.113, my emphasis), suggesting that contextualists don't indulge sceptical beliefs, even when reflecting on closure-based sceptical arguments. In the absence of any actual belief in such scepticism, it's highly unlikely that the reported intuition that it's true is just a disposition to believe it. Again, technically, contextualists failing to believe in scepticism isn't mutually exclusive with them having an extremely weak or systematically blocked disposition to believe it, but there's simply no reason to ascribe to them any such doxastic disposition, and every reason to doubt that they have one. To put it simply, while token dispositionalism is, at best, consistent with the evidence, it certainly isn't supported by it.

It could be objected that there's a plausible explanation for why some of the doxastic dispositions that constitute the Harman-Vogel paradox, specifically, are systematically blocked, namely, the fact that they have other intuitions working against them. Maybe, for instance, the intuition, or disposition to believe, that we don't know whether we have hands, even by ordinary epistemic standards, is blocked by the intuition, or disposition to believe, that our everyday knowledge

claims are true. After all, one, stronger disposition can obstruct another, weaker one. If, for example, I have a weak disposition to donate money to the needy, but a stronger disposition to purchase new video games, then the latter disposition will block the former; here, the former disposition to Φ would have manifested, if not for the fact that I'm also, and more strongly disposed to do something else, incompatible with Φ-ing. So, the argument might go, perhaps it's not so implausible to think that, in the case of sceptical arguments in particular, the intuitions in question are doxastic dispositions, some of which are systematically blocked by their stronger counterparts.

However, while this can make the idea that disbelieved contextualist intuitions are just systematically blocked doxastic dispositions seem like less of a stretch, it still doesn't give us any positive reason to adopt a dispositionalist thesis to begin with. We still have no tangible evidence that contextualists are disposed to believe (all) the propositions that they intuit, and thus no reason to treat the idea that their intuitions just are doxastic dispositions with more than a modicum of seriousness.

3.2.3 Familiar Worries

Token dispositionalism also succumbs to several of the objections that befell other, extant minimalist accounts of intuition. For one, dispositionalism likewise fails to respect the phenomenology of intuitions. To elaborate, recall that, if S would, reliably enough, Φ if *C*, then S has a disposition to Φ if *C* even when it isn't currently manifesting, to wit, when S isn't occurrently Φ-ing. I am, for example, at this very moment, disposed to take an umbrella out with me when it's raining, even though it's currently not, and that disposition isn't manifesting. After all, it's still true of me, right now, that I *would* take an umbrella out on such an occasion: I still harbour the tendency, or habit, to do so. Consequently, according to dispositionalism, S has the intuition that *p* even when that doxastic disposition isn't manifesting: they're still disposed to believe *p*, and thus still find it intuitive, even when they don't currently believe it. However, even supposing that some beliefs are conscious, merely *being disposed* to believe something, in itself, isn't a conscious mental episode. Right now, for example, I have a very strong disposition to believe that it's raining when it is and I'm outside, but there's no particular feeling that accompanies this fact. Dispositions, in general, are 'standing' mental states, waiting 'silently' in the background to be triggered by the right circumstances. Even when a disposition to experience a conscious event – such as a disposition to feel angry under certain conditions – does manifest, it's still the Φ-ing, not the disposition to Φ itself, that's consciously experienced. As such, it's difficult to reconcile dispositionalism with the aforementioned phenomenology of contextualists' intuitions. While

doxastic dispositions are unconscious, standing states, such intuitions, as we've already seen, are conscious mental events that their subjects 'feel'.

Indeed, this is a fairly common criticism of dispositionalism in general. Pust (2017), for example, observes that, if I'm not occurrently enjoying the right kind of phenomenology, then I'm not 'intuiting' anything, even if I'm disposed to believe very many things:

> "It is quite possible to have, at a time, a large number of dispositions to believe while failing to host, at that time, a single intuition" (§1.2).

Bengson (2014) makes a similar point here, noting that, if it doesn't *consciously* strike one that *p* is true, then one isn't intuiting *p*, even if they're disposed to believe it:

> "[O]ne's colleague might be disposed to think that contradictions cannot be true, but she is not right then and there intuiting this; as we say, it does not at that moment consciously *strike* her that contradictions cannot be true: the intuition is absent" (p.18, emphasis in original).

Admittedly, Bengson's project is very different from my own; he has in mind a particular conscious episode from the start, which, he says, some philosophers, including himself, call 'intuition'. His objective is merely to characterise this mental state, not to provide an account of intuition-talk *per se*, and certainly not the intuition-talk of contextualists specifically. Nonetheless, we can see that the same objection applies to the interpretive thesis that contextualists' intuitions are dispositions to believe.

Now, Sosa's dispositionalism arguably attempts to circumvent this phenomenological objection. Sosa claims that the intuition that *p* is a kind of disposition to believe *p* with particular manifestation conditions. Specifically, an intuition that *p* is a disposition to believe *p* that will manifest at time T1 if:

i) If at T1 S were merely to understand fully enough *p*, then S would believe *p*
ii) At T1, S does understand *p*
iii) *p* is abstract
iv) At T1, S thinks occurrently of *p* (Sosa, 1998, p.259)

Sosa's fourth criterion, specifically, could be interpreted as trying to subvert the phenomenological objection, by specifying that S only intuits *p* when they're occurrently (and, we can charitably assume, consciously) thinking about it – in other words, intuition is an occurrent, conscious mental episode, rather than an unconscious, standing inclination.

Does Sosa's version of dispositionalism – again, interpreted as a 'being disposed', not 'feeling tempted' account of intuition, for the time being – escape, at least, the phenomenological objection? I think not. Even when Sosa's four criteria aren't met, S still has the relevant disposition to believe *p* if they would do so (often enough) if they were. So, on his view, S still has the intuition that *p* even when they're not, for example, consciously thinking about *p*. That intuition, if it really is the disposition, itself, to believe *p* when i-iv are satisfied, is still just an unconscious, 'standing' mental state, waiting 'silently' in the background for those conditions to be met. In other words, rather than identifying intuitions with an occurrent, conscious mental event, Sosa's view still only identifies them with an unconscious disposition that manifests when a conscious mental event is happening. A conscious event is a manifestation condition of the intuition, but the intuition itself is still decidedly unconscious.

Sosa could reply that i-iv are just the conditions under which S has the intuition that *p*, rather than the manifestation conditions of that intuition. That way, S only intuits *p* when they're enjoying a conscious mental event with content *p*. However, then Sosa's view would only be that one is *disposed to intuit p* when i-iv are satisfied, and would no longer qualify as a 'dispositionalist' account of intuition, as I'm using the term, according to which intuitions are neither what dispose us to Φ, nor the Φ-ing that we're disposed to, but the disposition itself.

Moreover, even if Sosa's dispositionalism did survive the phenomenological objection, his account fails to capture contextualists' intuitions, specifically, in at least two other respects.

First, if the intuition that *p* were just a disposition to believe *p* with manifestation conditions i-iv, then, when S had the intuition that *p*, they would also have a disposition to believe *p* when those criteria are satisfied. However, contextualists can intuit *p* without harbouring such a doxastic disposition. For example, they find the claim that we don't know whether we have hands, even by ordinary epistemic standards, intuitive, at least when reflecting on AI, but have no disposition to believe it, even when they understand it 'fully enough', and are occurrently thinking about it. Thus, they can intuit *p* without possessing a doxastic disposition with Sosa's four manifestation conditions. Therefore, their intuition that *p* can't just be said disposition.

Second, Sosa's criterion iii entails that only abstract propositions can be intuitive. To clarify, if a manifestation condition of S' disposition to believe *p* is that the thing S is disposed to believe, *p*, is abstract, then the disposition only ever manifests in S believing abstract things. Since the disposition to Φ is just the disposition to do whatever Φ-ing constitutes, if Φ-ing always consists of believing abstract propositions, then the disposition to Φ is just the disposition to believe abstract things. So, because Sosa's view is that the intuition that *p* is a disposition to believe *p* that will only manifest if *p* is abstract, for him, an intuition is a disposition to believe an abstract

proposition, specifically[36]. I read an abstract proposition as one that doesn't specify any particulars, and that is, therefore, meant to generalise to all possible cases (cf. Ludwig, 2010, p.431). That describes the contents of at least some contextualist intuitions. The closure and sensitivity principles, for example, which are both reported by, contextualists, to be intuitive, are universal declarations about any and all cases of knowledge – any and all cases wherein a subject, S, knows a proposition, *p* – which fail to specify any particular subjects or propositions. However, there are also obvious counterexamples. In Cohen's discussion of the contextualist solution to sceptical paradoxes, for instance, he says:

> "[O]ur intuition that the sentence 'Smith knows his car is parked in lot 2' is true is explained by the fact that at ordinary contexts, the statement is in fact true. Our apparently conflicting intuition that the sentence 'Smith knows his car has not been stolen' is false results from our shifting to a stricter context when we consider the possibility that Smith's car has been stolen. At this new context, the sentence is indeed false" (2005, p.200).

In this excerpt, Cohen describes two, apparently conflicting intuitions. The first is that a particular subject, Smith, knows a particular proposition, that his car is parked in lot 2. The second is that the same subject doesn't know another, particular proposition, that his car has not been stolen. So, *contra* Sosa's dispositionalism, contextualists' intuitions aren't always abstract; sometimes, they specify particulars.

It might be objected that, when philosophers use names like 'Smith' in hypothetical vignettes, like the one above, they're meant to be read as placeholders for any epistemic subject, just as 'S' is meant to stand for any subject in formulations of CPK, for example. More generally, one might claim that our intuitions about such 'thought experiments' are actually about whether anyone would know any relevantly similar proposition under any relevantly similar conditions, and that they do, therefore, have abstract (or generalisable) content. However – aside from being a somewhat contrived way of making the theory that all intuitions are abstract fit the available data – this doesn't reflect actual philosophical practice. Philosophers are just as inclined to use intuition-talk, and its synonyms, when discussing *actual*, not hypothetical cases, where they certainly mean names to denote particular people, and not just to serve as placeholders, for

[36] Elsewhere, Sosa also says that the content of 'rational' intuitions, at least, is always "explicitly or implicitly modal (i.e. attributes necessity or possibility)" (2007c, p.101). In Chapter 4, in my discussion of the view that intuitions are 'intellectual seemings', I'll argue that contextualists' intuitions don't always have modal content either.

example. Cullison, for instance, has built a catalogue of real-life Gettier cases, and describes his reaction to them using the same kind of language that contextualists use to describe their own intuitions[37]. At the very least, the onus of proof is once again squarely upon my opponent to show that philosophers don't intend such language to be taken at face value; that intuitions are never really about such particulars, but about more general, abstract matters.

So, we have two, further reasons to reject Sosa's dispositionalism in any case. It's important to understand though that my arguments against such a dispositional account of *contextualists'* intuitions, specifically, aren't all direct criticisms of Sosa, since he explicitly suggests limiting his account of intuition-talk to certain sub-domains of philosophical enquiry, which may exclude contextualist literature. Specifically, he offers an account of the intuition-talk that's used in areas that deal predominantly in conceptual analysis, but says that "the use of intuitions in philosophy should not be tied exclusively to conceptual analysis" (2007c, p.100), because the vast majority of philosophical practice simply doesn't fit this description. Debates over utilitarianism versus deontology in ethics, or internalism versus externalism in epistemology, for example, are "disputes about something more objective than just a description or analysis of our individual or shared concepts of the relevant phenomena" (*Ibid.*); they're disagreements over the actual phenomena in question, such as knowledge or morality itself, and not just over what our, potentially inaccurate, concept of those things is. So, my argument here is just that Sosa's account, despite some superficial textual evidence in its favour, doesn't capture contextualists' intuitions, specifically. Whether it works better as an account of the intuition-talk that's being used in other areas of philosophy is beyond the remit of this project.

Moving on, another problem facing dispositionalism, apart from its inability to account for the phenomenology of contextualists' intuitions, is that doxastic dispositions are rationally criticisable. Specifically, as Sosa (2007a), for example, observes, a disposition to believe *p* may be too strong or too weak, relative to the strength of one's reasons to believe *p* (p.51). I'm criticisable if, for instance, I only have weak reason to believe *p*, but am strongly disposed to believe it, or if, conversely, I have strong reason to believe *p*, but am, at best, only weakly disposed to believe it. As I observed previously, contextualists' intuitions do, quite often, run contrary to their reasons for belief. Thus, interpreting their intuition-talk as referring to their doxastic dispositions, just like reading it as referring to their beliefs, would burden them with excessive irrationality, and be uncharitable.

[37] http://www.andrewcullison.com/2009/06/gettier-in-real-life/ [Accessed 31.08.2020]

Moreover, as with beliefs, doxastic dispositions are rationally criticisable, in part, because, normally, S can bring how disposed they are to believe *p* into line with their reasons to believe *p*, if they're so inclined. To elaborate, given that S can, as previously discussed, make whether and when they actually believe *p* sensitive to their reasons, and the strength of their disposition to believe *p* just is the measure or whether and when they actually believe *p*, the strength of that doxastic disposition can also be brought into alignment with their reasons, by extension. Again, this runs contrary to contextualists' intuitions, which are recalcitrant, and don't display such 'internal correctability'.

In conclusion, we shouldn't interpret contextualists' intuitions as doxastic dispositions. As a reminder, I'll address the possibility that contextualists' intuitions are doxastic *temptations* in the next chapter, and likewise conclude that they aren't.

3.3 Against Ordinary Language

The final minimalist reading of contextualists' intuitions that I want to refute is that, when contextualists call *p* intuitive, they just mean that people (especially laypeople) do or would (often enough) utter *p*. Under this interpretation, 'the argument from anti-scepticism' for EC appeals evidentially to empirical facts about how people do, or would, normally talk in, or about, real or imagined situations. For example, when contextualists say that, when we reflect on AI, we find 'scepticism' intuitive, they're just citing the observation that, usually, when reflecting on AI, people would (often enough) literally *say* that S doesn't know they have hands.

As with the other minimalist interpretations that I've discussed in this chapter, there's apt motivation to take this possibility seriously. It is, after all, quite natural to think that the evidence for contextualism ought to come from ordinary language. To elaborate, since EC is a linguistic thesis about the context-sensitivity of K-words, the best evidence in its favour would be if when and how competent English speakers used such terms in ordinary language displayed sensitivity to contextual parameters. So, it would only make sense if the evidence, or 'intuitions' cited in support of EC were facts about ordinary language. Indeed, DeRose even calls his seminal 2005 paper *the ordinary language* basis for contextualism, and, in an oft-quoted passage from *The Case for Contextualism* (2009), explicitly says that:

> "the best grounds for accepting contextualism come from how knowledge-attributing (and knowledge-denying) sentences are used in ordinary, non-philosophical talk" (p.47).

Also, while one way to interpret Cohen's attribution-talk is as referring to beliefs, as I previously discussed, another is to parse 'attributions' and 'denials' of knowledge as speech acts:

declarations that S does, or doesn't, know *p*, respectively. Thus, Cohen's contextualism might just rest on its ability to (charitably) explain common utterances of K-sentences (to reiterate, sentences of the forms 'S knows *p*' and 'S doesn't know *p*'). Indeed, in several places, Cohen talks as though intuitions are closely related to how people speak. For instance, at one point, Cohen is outlining the contextualist resolution to 'the lottery paradox', which is meant to be roughly analogous to the contextualist treatment of the Harman-Vogel paradox. The details of the former paradox, and its contextualist solution, aren't essential here. However, when discussing this paradox, Cohen will sometimes describe one of its constitutive data points as the 'intuition' that S knows they will lose (or have lost) the lottery (2005, p.200), while, at other times, he'll instead describe it thusly:

> "In both of these cases we are *inclined to say* that S does know that he loses" (1988, p.92, my emphasis).

So, what is, in one place, an intuition, is, in another place, a fact about what we're inclined to say. Also, when discussing the intuitiveness of 'scepticism', Cohen describes it as a pull toward *saying* that it's true:

> "[W]hen in the grips of a sceptical argument, we feel the strong intuitive pull of *saying* that our knowledge ascriptions are mistaken" (*ibid.*, p.302, my emphasis).

Finally, when presenting his contextualism as the most damage-controlling explanation for our paradoxical intuitions, Cohen says that it does justice to, among other things, what we 'say' we know in everyday contexts:

> "By supposing that the truth-value of an attribution of knowledge is context sensitive in just this way, we do justice to both the undeniable force of skeptical arguments and our strong inclination to *say* that we know things in everyday life" (*ibid.*, p.97, my emphasis).

As with attribution-talk, one way to read 'what we say' is as referring to beliefs. Sometimes, for example, I might describe myself as believing *p* by saying things like 'I would say *p*'. However, we could also take this terminology at face value: as making reference to what sentences we, literally, do or would utter.

However, while some arguments for EC might rely on empirical facts about ordinary language, 'the argument from anti-scepticism' doesn't.

For one, we can assume that most utterances are sincere, not deceptive. However, while believing and uttering *p* are distinct phenomena, since one might, for instance, surreptitiously believe *p* without ever actually saying it aloud, a sincere utterance of *p* requires a corresponding

belief that *p* – otherwise, it isn't sincere – and thus entails said belief. Therefore, if this interpretation, according to which people's intuitions are just their ordinary linguistic utterances, were true, we would expect subjects to believe their intuitions: we would expect such utterances to, on the whole, be sincere, and thus entail corresponding beliefs. But, again, contextualists' intuitions are, quite often, *not* believed by their subjects. So, it's an indirect consequence of rejecting doxasticism that we shouldn't understand said intuitions in terms of what people do or would *literally* say.

Indeed, if we look more closely at the passages quoted above, we see that Cohen doesn't straightforwardly refer to speech-acts in themselves, but rather to what we're 'inclined' to say, what we feel a strong 'intuitive pull' to say, and what we have a 'strong inclination' to say, respectively. Since we've already seen that contextualists' can intuit *p* without being disposed to believe *p*, and given that a disposition to sincerely assert *p* entails a corresponding disposition to believe *p* under the same conditions, since, again, sincere utterance entails belief, we shouldn't read this as referring to linguistic dispositions. However, we could instead interpret it as referring to conscious linguistic temptations: to *feeling* inclined, or 'pulled', toward asserting something, without necessarily actually uttering it, or even being disposed to (often enough) do so. Now, again, I'll discuss the relationship between contextualists' intuitions and temptations in the next chapter, where I conclude that aren't reducible to those either. However, I'll argue that that can, and often do, *ground* (cause, or explain) such temptations. For instance, I'll make the case that the relevant intuitions are a phenomenon that tempts, or encourages, but doesn't entail, assertion (because it encourages, but likewise doesn't entail, belief). That will explain the apparent connection in contextualist literature between intuitions and what people feel compelled to say, but don't, or wouldn't, necessarily, actually assert.

I'll briefly note that the ordinary language interpretation of contextualists' intuitions also succumbs to familiar worries. For instance, neither speech acts, nor empirical facts about them, are conscious mental episodes. However, I don't wish to overly labour these points.

In conclusion, we shouldn't interpret contextualists' intuition-talk as referring to ordinary language.

3.4 Chapter 3 Conclusion

In conclusion, contextualists' intuitions aren't outright beliefs, partial beliefs, dispositions to believe, or ordinary speech acts. That contextualists sometimes describe them as though they are is a further reminder of the need for a rational reconstruction of their intuition-talk in clearer, more coherent terms. Nonetheless, the tendency for contextualists to conflate intuitions with

such things suggests that intuitions are at least related, and perhaps familiar, to these various phenomena, and any account of them should respect this fact.

This concludes the destructive portion of my investigation into research question B. In the next 2 chapters, I undertake my constructive project, arguing that contextualists' intuitions are 'intellectual seemings', which are, in turn, a sub-category of 'epistemic feelings'. Finally, in Chapter 6, I consider some of the implications of this discovery.

Chapter 4 Intuitions as Intellectual Seemings

In the last chapter, in my destructive project, I showed that contextualists' intuitions are neither outright beliefs, partial beliefs, or dispositions to believe, nor ordinary speech acts. In this chapter, I begin my constructive project, arguing that contextualists use intuition-talk to refer – loosely, and perhaps unknowingly – to a phenomenon that has been variously labelled 'intellectual seemings' (Brogaard, 2013c), 'intellectual appearances' (Huemer, 2007, p.40), or 'apparent rational insights' (BonJour, 1998, p.114).

I'll argue for the position that contextualists' intuitions are intellectual seemings (as I'll elect to call them) in two stages. First, in §4.1, I'll demonstrate that contextualists use intuition-talk interchangeably with seeming-talk, outline the different kinds of seeming to which they might be referring, and argue that contextualists' intuitions are, specifically, of the perceptual (or phenomenological) variety. To this end, I'll employ both negative and positive strategies, that is, I'll argue both that the relevant intuitions *aren't* other kinds of seeming, and that they *are* perceptual ones. I'll achieve the latter step by appealing to numerous structural and metaphysical resemblances between the relevant intuitions and other, more paradigmatic classes of perceptual seemings, such as visual experiences. Then, in §4.2, I'll argue that the aetiology of contextualists' intuitions is specifically intellectual, rather than experiential, memorial, or introspective, and therefore conclude, overall, that they're *intellectual* perceptual seemings, specifically, and thus a variety of 'intellectual seemings' *simpliciter*.

In §4.3, I'll argue that no extant intellectual seemings account of philosophical intuitions in general successfully characterises contextualists' intuitions specifically. More exactly, these prior proposals introduce further conditions for being an intuition that exclude contextualists' intuitions, and thus undergenerate. I'll therefore adopt a novel intellectual seemings theory that foregoes these additional requirements.

In §4.4, I'll argue that my constructive thesis respects the evident connection and similarity between contextualists' intuitions and the various other phenomena with which they often conflate them, namely, outright beliefs, credences, dispositions, and ordinary speech acts. Thereby, my account helps to explain the temptation for contextualists to interchange talk of intuitions with talk of such things, and their propensity to confuse these different entities.

In sections 4.5 and 4.6, I'll develop the account further. First, in §4.5, I'll argue more comprehensively for the view that contextualists' intuitions, or intellectual perceptual seemings, are 'presentational' mental states. In other words, I'll argue that an intellectual seemings account

of intuitions can and should be reconciled with a presentational (or 'quasi-perceptual') account, despite existing concerns to the contrary. I'll thereby develop and defend a novel, hybrid theory of philosophical intuition. Then, in §4.6, I'll examine the phenomenology of the relevant intellectual seemings more closely, and argue that they aren't action-oriented doxastic temptations, *contra* Sosa, but rather proposition-oriented appearances of truth.

Finally, I'll conclude that contextualists' intuitions, at least, are *intellectual perceptual seemings*, that such seemings are *presentational* states, and that they're not doxastic temptations.

Later, in Chapter 5, I'll elaborate still further upon this position, by arguing that the relevant intellectual seemings aren't a *sui generis* phenomenon, as many are wont to assume, but rather a sub-class of what psychologists call 'epistemic emotions', and highlight some of the important implications of this discovery in Chapter 6.

4.1 Intuitions as Perceptual Seemings

In this section, I argue that contextualists' intuitions, at least, are perceptual seemings.

Contextualists often use intuition-talk interchangeably with seeming-talk. For example, at times, DeRose claims that his EC can explain what he explicitly calls the 'intuitiveness' of 'ignorance', and of the sensitivity principle (as we've already seen). Elsewhere, however, he says that his theory can explain, specifically, why these things 'seem' true, as he does here:

> "[M]y solution explains why that premise [ignorance] *seems* true and, more generally, why sensitivity *seems* necessary for knowledge" (2017, p.20, my emphasis).

Moreover, in this next passage, DeRose explicitly states that to explain (away) the 'intuitive power' of general scepticism is to explain (away) why it can 'seem' true:

> "[M]y response is to provide an explanation of how the skeptic's argument can make it *seem*… that we don't know-to explain (away) the *intuitive power* of the skeptic's case" (*ibid.*, p.115, my emphasis).

Both of these excerpts strongly suggest that, for DeRose, to explain intuitions is to explain why things seem true and, therefore, that intuitions are, for him, said seemings.

Seeming-talk is comparably popular in Cohen, and likewise used interchangeably with intuition-talk[38]. For instance, in this next passage, Cohen (2005) describes the fact that, when we consider AI, it 'seems' that our ordinary claims to knowledge are mistaken as a 'fact about' our intuitions:

> "But here is *a fact about those intuitions*. When we are in a sceptical frame of mind, it *seems* to us that we have made a discovery - there are lots of things we think we know… that in fact we do not know" (p.207, my emphasis).

This strongly suggests that, for Cohen, our intuitions are just what seems true to us and, therefore, that intuitions are, likewise for him, said seemings.

So, for both Cohen and DeRose, intuitions are seemings. However, seeming-talk can be used in a variety of different ways, to refer to several different phenomena. In the next sub-section, I disambiguate those uses, and pinpoint the specific species of seeming to which the contextualists are alluding, namely *perceptual* seemings[39].

4.1.1 Varieties of Seeming

The first way to use seeming-talk is what Chisholm (1957, p.45) called 'comparatively'. Speakers use seeming-talk comparatively when they say 'A seems like B', for example, and just mean that A bears some similarities to B. If, for instance, Carla is looking particularly pale today, then one might say 'Carla seems like a ghost', in reference to the visual resemblances between them. Of course, this use of seems need not entail, or express, that the speaker believes that A *is* B. For instance, asserting that Carla seems like a ghost in certain respects need not express, or entail, the belief that Carla *is* a ghost. Indeed, this kind of seeming-talk commonly appears in statements of the form 'A seems like B, but it isn't'. For example, a speaker might convey the fact that a certain bird looks like, or is otherwise very similar to, a swallow, but isn't, by saying 'that bird *seems* like a swallow, but it isn't'.

However, this first use of seeming-talk, fairly obviously, doesn't reflect how it's employed in contextualist literature. To elaborate, as Brogaard (2013a, p.214) observes, only if, when one says that something seems a certain way, they compare it with something else, is the use of 'seems' comparative. However, as we see in the excerpts above, contextualists don't use seeming-talk to compare things. They assert that things – like 'ignorance', or the sensitivity principle – seem true, not that they seem 'like' something else.

---

[38] See also, for example, Cohen (2005, p.199-200, 1988, p.92).

[39] I hope that this process of disambiguation will also assist my reader to locate the relevant species of seeming in their own mental lives, in case the phenomenon is, at first, 'unintuitive'.

Next, I'll discuss the epistemic uses of seeming-talk, of which there are two.

The first epistemic use of seeming-talk is when a speaker uses 'seems' to hedge their claims. When, instead of merely asserting *p*, S says, or writes, 'it *seems* that *p*', they might still mean to claim that *p*, but only tentatively: the addition of 'seems' conveys that they're not strongly committed to *p*, that their confidence in *p* is only just high enough for outright belief, that their view is liable to change easily in the future, that they're receptive to alternatives, and so on. In other words, 'seems' can be used to hedge in the 'shielding', as opposed to 'approximating', manner, a distinction first made by Prince *et al.* (1982). To elaborate, approximating hedges are when people say things like 'roughly, *p*', or 'about *p*', and their purpose is to convey that the speaker is committed, not to *p* specifically, but to something in its general vicinity. Meanwhile, shielding hedges are when people say things like '*p*, I think', or '*p*, I believe', and their purpose is to convey reduced confidence in *p*, and thus guard the speaker against excessive rational criticism if they're mistaken: because they aren't strongly committed to *p*, and are open to alternatives, it doesn't reflect as poorly upon them if *p* is false. This corresponds to what Chisholm (1957, p.44) simply called *the* epistemic use of 'seems', wherein the speaker "hedges, giving [themselves] an out" (1982, p.139).

Andow (2017) and Cappelen (2012, p.41) have both argued that at least many uses of intuition-talk in contemporary analytic philosophy should be read as shielding hedges. The former, for example, cites the correlation between the increased use of intuition-talk among philosophers (2015b) and the more general increase in hedging across academia (2015a) as evidence that the former trend reflects the latter. If philosophers were using seeming-talk to hedge, and also using it interchangeably with their intuition-talk, then that would be good reason to take such a proposal seriously.

However, we at least shouldn't read contextualists' seeming-talk, specifically, in this way either, for two reasons.

First, according to this interpretation, contextualists use seeming-talk to hedge their *beliefs*. When they say that 'ignorance' seems true, for example, they just mean that they believe 'ignorance' is true, albeit tentatively. However, as we've seen, propositions that 'seem true' to the contextualists are also ones that they intuit and, as I argued in Chapter 3, they don't always believe their intuitions. So, something can seem true to them without them believing it. Consequently, things seeming true to them isn't just them believing those things, even hesitantly. Moreover, even if we overlook how their intuition-talk is used, their seemings, themselves, aren't always things that they outright believe. In one of the excerpts above, for instance, Cohen says that, when he reflects on AI, it 'seems' to him that our everyday knowledge claims are mistaken,

which is certainly not something that he outright believes, even tentatively. Thus, again, contextualists don't simply use seeming-talk to hedge, that is, to express weak outright beliefs.

Second, shielding hedges convey that the speaker is less confident in a claim than it might otherwise appear if they expressed it straightforwardly. However, as we've seen, contextualists' seemings are synonymous with their intuitions, and they take intuitions to directly support certain claims. Indeed, even if we overlook how their intuitions are used, again their seemings themselves are sometimes cited as a fact in favour of believing something. The sensitivity principle, for example, is something that 'seems' true to DeRose, and, in some places, it's precisely such a 'seeming', not explicitly an 'intuition', that he says speaks in favour of believing it (see, for instance, 2017, p.19-28). Either way, when contextualists say that something seems true, they don't mean to express reduced confidence in it, but to indicate a fact that increases or grounds their confidence. Therefore, again, contextualists' seeming-talk isn't being used to hedge.

The second epistemic use of seeming-talk is to convey that something is "subjectively probable" (Brogaard, 2013a, p.214). To elaborate, when S says 'it seems that *p*', for example, they could mean that *p* is (most, or sufficiently) likely, conditional on the (total, available, relevant, or salient) evidence. If, for instance, a homicide detective, having collated evidence suggesting that Smith is the murderer, says "it seems that Smith is the murderer", they could mean that, on the basis of the evidence they've collected thus far, Smith is the most, or reasonably, likely to be guilty. Unlike with the hedging use of seeming-talk, what 'seems' true to S, in this sense, can differ form what they outright believe, because, even when S deems *p* probable enough, conditional on the evidence, to say that it 'seems' true, they still might not be confident enough to actually believe it.

However, we shouldn't read contextualists' seeming-talk as expressing subjective probabilities either. To elaborate, if this interpretation were true, then what seems true to contextualists should correlate with their credences, that is, how likely they find things, conditional on their reasons/evidence. Specifically, only when they find *p* most or sufficiently likely should they declare that *p* 'seems' true. However, as we've seen, something seeming true to contextualists is the same as them intuiting it and, as I've already demonstrated, in Chapter 3, contextualists can intuit *p* without finding *p* most, or even reasonably, likely, on the basis of their reasons/evidence. Therefore, when contextualists say that *p* seems true, they're not just expressing subjective probabilities.

With the comparative and epistemic uses of seeming-talk out of the way, that just leaves the 'perceptual' use (Brogaard, 2013a) (otherwise known, in Jackson (1977), for example, as the 'phenomenal' use). To elaborate, sometimes, when S says 'it seems that *p*', they mean to describe how the world phenomenologically appears (or how it, in some, broad sense, 'feels'). This species

of seeming is best explained by way of example. Consider, for instance, the infamous Müller-Lyer perceptual illusion (*Fig 1*).

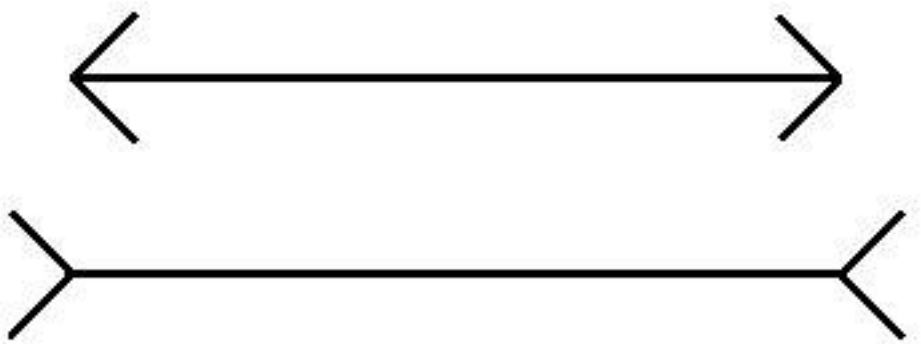

Fig 1 Müller-Lyer Illusion

In the above image, the bottom line, in some sense, 'seems' longer than the top one, even though it isn't. However, it will continue to seem that way, in this respect, even if you believe, have overwhelming evidence to believe, and are extremely confident that, it isn't. Therefore, when S says 'the bottom line seems longer', they're not just reporting an outright or partial belief, hedged or otherwise, or a subjective probability. Nor are they comparing the bottom line to something else; how long the bottom line seems, in this sense, isn't just a fact about how long it is relative to something else (say, the top line), but a property of the bottom line in itself: how long *it* seems. Rather, this sense of 'seems' refers to how the line *phenomenally* appears or feels to S (in this case, how it, quite literally, looks)[40]. Therefore, some seeming-talk expresses the subject's phenomenology, or, as I'll call it, their 'perceptual seemings'. As Bealer (1998) says:

> "Here 'seems' is understood, not in its use as a cautionary or "hedging" term, but in its use as a term for a genuine kind of conscious episode" (p.207)[41].

We can think of the subject's visual experiences – such as the look of the lines in *Fig 1* – as paradigmatic examples of these perceptual (or phenomenal) seemings.

While intuitions are quite different from visual experiences in a number of ways (I'll say more on that later), I argue that contextualists' intuitions, like visual experiences, are also perceptual seemings. Apart from being the only candidate remaining for the referent of their seeming-talk, I argue that contextualists' intuitions are perceptual seemings on the grounds that the two

[40] Perceptual seemings can obviously ground partial or outright beliefs, hedged or otherwise. S might, for example, believe, or be somewhat confident, that the bottom line in *Fig 1* is longer than the top line because it phenomenally seems so. Consequently, it seeming to S in one way can ground it seeming to S in another.

[41] In this passage, Bealer is talking specifically about intuitions, while I'm, for now, talking about perceptual seemings more generally.

phenomena display numerous similarities. Specifically, there are ten respects in which contextualists' intuitions resemble perceptual seemings in general, including sensory, such as visual, seemings.

4.1.2 Resemblances Between Intuitions and Perceptual Seemings

Contentful. Contextualists' intuitions and perceptual seemings both have content.

For example, when it visually seems to S that there's an apple on the table, it seems to S that *something*, namely, that there's an apple on the table. That something is the 'content' of S' perceptual seeming, or what the state is 'about', to wit, what it is that seems. Content can be used to differentiate specific perceptual seemings. For instance, it seeming to S that there's an apple on the table, and it seeming to S that there's a banana on the table, are different states, because, while both are cases of perceptual seeming, they have different content, that is, different things 'seem' to S. If perceptual seemings weren't contentful, it would be much harder to delineate such experiences. Likewise, when contextualists intuit that, for example, 'ignorance' is true, they intuit *something*, namely 'ignorance'. That proposition is, again, the content of their intuition, or what the state is 'about', to wit, what it is that's intuited/intuitive. The intuition that the sensitivity principle is true is a different intuition because, while both are intuitions, they have different content, that is, different things are being intuited.

Representational. Contextualists' intuitions and perceptual seemings are both representational.

A state is representational if it represents the world as being a certain way, namely the way it would be if its content were true (Peacocke, 2004, p.99). We can test for this feature by looking at whether the state in question can be correct or incorrect, relative to the world. Because nonrepresentational states are silent about the way the world is – because they don't represent it as being any particular way – they're neither correct nor incorrect about the world[42]. Hopes, for example, plausibly don't represent the world as being any particular way. The hope that it will be sunny tomorrow, for instance, has content, namely 'it will be sunny tomorrow', or what is hoped for. However, a hope is silent on the matter of whether its content is actually true: it doesn't represent the world as being, or not being, the way it's hoped to be. *Qua* hoping that it will be

[42] Terms like 'correct' and 'incorrect' are more appropriate here than 'true' or 'false', for two reasons, First, because, strictly speaking, propositions, not mental states, are truth-value bearers. For example, it only makes sense to say 'S' belief is false' because 'belief' is ambiguous between the mental state itself and its propositional content; here, what we really mean is '*p* is false', not that the attitude, believing-*p*, is. A representation with false propositional content is not a 'false' state *per se*, but one that's 'wrong' or 'inaccurate' in the sense that it misreports the world. Second, because, as I explain in footnote 43, representations need not have propositional, or truth-apt content at all.

sunny tomorrow, S *merely* hopes that it will be sunny tomorrow, and is neither given, nor under, the impression that it will be (c.f. Bengson, 2015). Accordingly, hopes are neither correct nor incorrect, relative to the world. For example, a hope isn't 'wrong' just because its content is false. By contrast, representational states specify 'accuracy conditions', or how the world must be for them to be correct (Tye, 2009, p.252; Siegel, 2010; Searle, 1983, p.43; Peacocke, 1983, p.5; Crane, 1992, p.139). Put simply, because the state represents the world as being a certain way, the world must really be the way it's represented as being in the state, or else the state is inaccurate. The most paradigmatic example of representational states is beliefs. The belief that it will be sunny tomorrow, for instance, likewise has content, namely 'it will be sunny tomorrow', or what is believed. Here though, the state depicts the world as being a certain way, namely, as a world where it will be sunny tomorrow: believing is a 'representational attitude'. Consequently, if it won't be sunny tomorrow, then the belief is incorrect (it *mis*represents the world). Generally then, if a state can be deemed correct or incorrect, relative to the world, then it's representational, and *vice versa*.

Perceptual seemings are representational, in the broad sense discussed here, because they too can be correct or incorrect, relative to the world. In *Fig 1*, for example, the bottom line perceptually seems longer then the top one, but it isn't. This is to be contrasted with, say, a perceptual seeming as-of an apple on the table, when there really is one. The former is an in accurate representation of the world, the latter an accurate one[43]. Contextualists' intuitions are likewise representational. They have content, and that content can be correct or incorrect, relative to the world. Recall that contextualists deem at least some of their intuitions, like the intuition that the sensitivity principle is true, to be mistaken, and others, like the intuition that 'closure' is true, to be correct. If they're using intuition-talk to refer to something that can be right or wrong, relative to how things really are, then they're referring to, specifically, representational states.

Presentational. Contextualists' intuitions and perceptual seemings are both presentational.

A sub-category of representational states are presentational states, which, more than simply representing the world as being a certain way, also *present* the world, to their subject, as being

[43] It isn't a precondition of being representational, in this broad sense, that a state has propositional content, meaning content with a truth-value. Representational states merely need to have content that can, in some sense, align with reality, or fail to. A portrait, for example, could be said to more or less accurately represent its subject, but we need not construe paintings as having propositional content. Thus, it isn't a problem for me if some perceptual seemings, like visual experiences, have non-propositional content; the fact that reality can align with (be similar to, for instance) how it's 'pictured' in one's experiences, or fail to, is all I need for the purposes of this comparison.

that way (Bengson, 2015, p.716). To elaborate, presentational states are how the world is 'given' to S, rather than how S 'takes' the world to be. They're essentially the agent's cognitive 'inputs' – the way the world is initially represented to them as being, and the basis on which they decree the world to be such and such a way – rather than their cognitive 'outputs', or how they've actually come to believe the world is. To help clarify the distinction between presentations and representations *simpliciter*, consider, for example, the difference between beliefs and experiences. Beliefs are 'cognitive endorsements' (*Ibid.*, p.717): in outright believing *p*, S is committed to the world being a certain way, namely a *p*-world. However, their belief that *p* doesn't 'say' to S that *p* is true. A belief is an 'output' representation, or how S has decided the world to be, rather than an 'input' representation, or how the world is 'given' to them. By contrast, experiences represent the world being a certain way, as previously discussed, but also present it, to their subject, as being that way. For example, a visual experience as-of there being an apple on the table 'tells' S, we might say, what the world is like, namely, that there's an apple on the table. That data then informs, among other things, the formation of S' beliefs[44].

Because presentations are how the world is given to S, not how they take the world to be, it being presented to S that *p* doesn't entail that S (outright or partially) believes *p*. Indeed, while an agent usually takes the world to be the way it's presented to them (and often, but not invariably, *because* it is thus presented to them), they might not: one can refuse to acquiesce to how the world is presented to them, and fail to form any corresponding belief, if they suspect their presentations to be misleading. In Bengson's terms, S can *have* the impression that *p* – have it presented to them as though *p*, by, say, an experience – without being *under* the impression that *p*, meaning without believing *p* (*Ibid.*). Or, as I might put it, the world can be *given* to S a certain way, without them *taking* it to be that way. By contrast, because nonpresentational representations are how S takes the world to be, it's not possible for them to have a nonpresentational representation with content *p* without outright or partially believing *p*: one state-type tautologically entails the other. Accordingly, a reliable smell test for whether a state is presentational is whether it's representational but non-committal. I've already defined representational states above and, by non-committal, I just mean not belief-entailing (to wit, it's *possible* to have the state without, outright or partially, believing its content). If a state is representational but committal, then it's an output representation, or how S takes the world to

[44] Beliefs can also, of course, be used as inputs for S' cognitive mechanisms, and inform the creation of further beliefs, and other, 'downstream' mental states, but they are simultaneously 'outputs', being conclusions previously drawn by those mechanisms. This is also not to deny that some believed propositions might, simultaneously, be being presented to S. The point is just that the contents of beliefs aren't, *qua* contents of beliefs, presented as being true, while the contents of experiences are.

be. If a state is representational but non-committal, then it's an input representation, or how the world is 'given' or 'presented' to S.

Perceptual seemings are presentational, because they pass this smell test with flying colours. For example, it perceptually seeming to S that, in *Fig 1*, the bottom line is longer is clearly an 'input' rather than an 'output' representation, because, as previously stated, the lines will continue to phenomenologically 'seem' this way to S even if they believe they aren't. Indeed, it's near-consensus that it can perceptually seem to S that *p* without S actually believing *p* (and that this extends to perceptual seemings with an intellectual, rather than experiential, aetiology) (Bealer, 1999, p.31; Huemer, 2001, p.99-100; Sosa, 2006). If it can perceptually seem to S that *p* without S believing *p*, then perceptual seemings are non-committal representations. Thus, the phenomenal appearance of the lines in *Fig 1* is how the world is given to S, not necessarily how they take it to be. Contextualists' intuitions likewise pass the smell test for being presentational. As I've already argued, they represent the world, and we've already seen that they don't entail belief in their content: a contextualist can intuit *p* without believing *p*. Again then, they're how the world is given to the contextualists, not necessarily how they take things to be. I'll argue more fully for this later, in §4.5, wherein I discuss the presentationality of intuitions in greater detail.

Non-Factive. Contextualists' intuitions and perceptual seemings can both be mistaken.

The fact that representational states are states that can be correct or incorrect, relative to the world, isn't mutually exclusive with some representational states being infallible. For a state to be representational, it just has to be the case that, in principle, *if* the world were different from how it's 'pictured' in the state, then the state *would* be mistaken, which is compatible with there being some representational states that, in actuality, never fail to report how things really are. However, it can perceptually seem to S that something is the case even if it isn't and, therefore, perceptual seemings are fallible, or non-factive[45], representational states. For instance, it can perceptually seem to S that the lines in *Fig 1* are different lengths, when they aren't. Likewise, as we've seen, contextualists often take their intuitions to be mistaken, and their intuition-talk, therefore, means to refer to fallible representations, specifically.

Phenomenal. Contextualists' intuitions and perceptual seemings are both phenomenal.

Perceptual seemings are, by definition, phenomenal. Again, for it to perceptually seem to S that *p* is just for it to *phenomenologically* appear to S that *p*, and thus, it isn't possible for it to

[45] This references the idea that you can't validly infer *p* from the fact that it perceptually seems to S that *p* (things aren't always what or how they seem). Contrast this with knowledge, which is 'factive', meaning, one can validly infer *p* from the fact that S knows *p* (everything known is a fact).

perceptually seem to S that *p* without S enjoying some form of phenomenology. For example, when it perceptually seems to S that the bottom line in *Fig 1* is longer than the top one, S enjoys some visual phenomenology: the colour and shape of the lines, against a white background. Likewise, when contextualists' intuit that, say, 'closure' is true, they experience phenomenology. After all, we've already seen that contextualists frequently describe intuitions as conscious (or phenomenal) mental events that they 'feel'. This strongly suggests that, when contextualists say that *p* is intuitive, they just mean that is 'feels', or *phenomenologically seems* to them, that *p*. Later, in §4.6, I'll examine the specifics of that feeling more closely.

Gradable. Contextualists' intuitions and perceptual seemings both come in degrees of vividity, or exactness.

Perceptual seemings come in degrees of exactness, or vividity. Consider, for example, a visual experience as-of there being an apple on the table. If said experience is occurring under ideal lighting conditions, and nothing is inhibiting S' visual faculties, then what perceptually seems to S will be something *very specific*: that there's an object on the table that's distinctly an apple, and an apple of a particular shape, size, and colour. On the other hand, however, if S is short-sighted, and removes their glasses, then what perceptually seems to them to be on the table will be something relatively less exact, and increasing the distance between them and the object will only exacerbate the imprecision. In such cases, it might only perceptually seem to S that *something* is on the table, of *roughly* a certain shape and size. In other words, the content of perceptual seemings can be more or less vivid.

Contextualists' intuitions are also gradable in this respect. For one, recall that DeRose describes his intuition that 'closure' is true as being much clearer than his, comparatively vague intuitions toward 'ignorance'. While 'closure' feels to him, very distinctly and precisely, *true*, his feelings toward whether S knows they're not a BIV are much hazier and more ambiguous: the latter feels, roughly, somewhere between true and false. In other words, the content of contextualists' intuitions can likewise be more or less vivid.

Recalcitrant. Contextualists' intuitions and perceptual seemings are both recalcitrant.

As I've already touched upon, even if S takes themselves to have overwhelming evidence that things aren't the way they perceptually seem to be, things will (typically) continue to perceptually seem that way to them. For example, even when S has measured the lines in *Fig 1* several times, and takes themselves to therefore have overwhelming evidence that they're the same length, the bottom line will continue to perceptually seem longer. Perceptual seemings aren't, therefore, sensitive to evidence. This is, at least partly, because S can't exert control over how things

perceptually seem to them, and thus no amount of effortful reflection on S' part can force their perceptual seemings into agreement with (the rest of) their evidential *corpus*. For instance, no matter how hard they look at, or think about, *Fig 1*, they can't make the bottom line stop looking longer. In other words, perceptual seemings don't display what I've previously called 'internal correctability'. They are, instead, recalcitrant. Once again, this is the same for contextualists' intuitions, as we've already seen. Neither phenomenon is sensitive to evidence, partly because neither is internally correctible.

Involuntary. Contextualists' intuitions and perceptual seemings are both involuntary.

Perceptual seemings are, moreover, involuntary. Just as it doesn't stop perceptually seeming to S that *p*, even when S has overwhelming evidence, and believes, that *p* is false, partly because perceptual seemings aren't 'internally correctible', S also cannot prevent the world from perceptually seeming a certain way to them in the first place. By contrast, the world will perceptually seem a certain way to S whether they want or intend it to or not, and whether they think that it's really that way or not. In other words, perceptual seemings aren't something that S voluntarily makes or even allows to happen, but something that involuntarily happens *to* S, without their say. As Bealer (2015) puts it, a presentational state "comes upon us unbidden, as if from without" (p.721). For example, just as no degree of internal reflection on S' part can correct the way that the lines in *Fig 1* perceptually seem to them, S also cannot stop them from seeming to be of different lengths from the start, nor does said seeming require anything like S' permission, intention, or agreement to appear. Another way to put this is that, just as perceptual seemings aren't *subject* to S' agency (they're not internally correctible, for instance), they're also not the *result of* S' agency.

We have every reason to think that contextualists' intuitions are likewise involuntary in this respect. Again, contextualists sometimes don't endorse the intuitions that come upon them. As I've previously argued, this shows that said intuitions are recalcitrant, and not internally correctible. However, it also suggests that contextualists can't prevent their intuitions from occurring. For we would expect contextualists to stop themselves from intuiting something that they take to be false as true in the first place, if they could, again if only to spare themselves the cognitive dissonance. Thus, if they continue to experience the feeling that, say, the sensitivity principle is true, despite not believing it, then said intuition can't be prevented. At the same time, the fact that said intuition appears even when contextualists don't intend it to or agree with it shows that it's an involuntary state: something that just happens to them, without their say.

Non-inferential. Neither contextualists' intuitions, nor perceptual seemings, are consciously inferred from something else.

Although preceding states, at the subpersonal level, might cause a perceptual seeming to arise, such seemings aren't formed on the basis of any such states, that is, inferred from them, or facts about them (cf. *Ibid.*). After all, such seemings are, to reiterate, paradigmatically presentational states, and thus not cognitive 'outputs', or how S takes things to be, on the basis of something else, but how things are 'given' to S. For example, while subpersonal processes (such as visual mechanisms) might causally explain perceptual experiences (such as visual experiences), agents don't consciously infer their perceptual experiences from those processes, or facts about them. S doesn't, for instance, consciously infer that the lines in *Fig 1* are of different lengths from any preceding premise or consideration (indeed, we might add that S explicitly infers that they are not, from the fact that they have measured them multiple times). Rather, the lines just perceptually seem that way to them, on no inferential grounds.

It's near-consensus that philosophical intuitions also aren't consciously inferred from other states (Bealer, 1996, 2000; Pust, 2000; Brogaard, 2013b, 2013c). As Rini (2015) says, "it is… widely accepted that intuitions are 'non-inferential'" (p.3). Gopnik and Schwitzgebel (1998), for example, describe intuitive judgements as being:

> "not made on the basis of some kind of explicit reasoning process that a person can consciously observe. Intuitions are judgments that grow, rather, out of an underground process, of whatever kind, that cannot be directly observed" (p.77).

This suggests that intuitions are caused by subpersonal (or 'underground') processes, but not consciously reasoned or inferred from anything. Deutsch (2010) has a similar view, wherein to say '*p* is intuitive' is just to say that one can know *p* without consciously inferring it from anything else, and 'intuiting' just means 'knowing directly'. The view that intuitions are non-inferential also has a strong historical precedent. Turning to the classic 'moral intuitionists', Price (1758), for example, believed that self-evident propositions can be known purely by reflecting on their content, a process called 'intuiting' them, resulting in a kind of direct knowledge called an 'intuition'. He defined 'intuiting' as coming to know something "without making any use of any process of reasoning" (*Ibid.*, p.159), and we can assume that 'reasoning' includes, for example, making conscious inferences.

Contextualists' intuitions, specifically, appear to be likewise non-inferential, because contextualists' arguments 'bottom out' at said intuitions. For example, the support for 'closure' is simply its intuitiveness, and no further defence of said intuition is offered. Indeed, we see this pattern time and again in the relevant literature: claims are substantiated by intuitions, but the argumentative regress terminates there. The fact that intuitions themselves are never supported by other premises suggests that intuitions aren't, in turn, inferred from any such epistemically

prior considerations. Put simply, one would expect that, if contextualists consciously inferred their intuitions from other premises, then they would know the basis of said intuitions, and that, were such further support known to them, then they would cite it in the text, since that would provide more epistemically foundational evidence for their thesis, and also save them from having to base their views solely on this hazy thing called 'intuition', thereby strengthening their argument. So, we can safely say that contextualists' intuitions mark the beginning of an inferential chain, and aren't, themselves, inferred from anything more foundational.

Tempting. Contextualists' intuitions and perceptual seemings can both consciously tempt S to believe their content, especially if they don't already.

Normally, it perceptually seeming to S that *p* causes S to believe *p*. However, the transition from perceptual seeming to belief isn't typically accompanied by, or mediated via, some conscious temptation to form said belief. Usually, S simply experiences the perceptual seeming and spontaneously believes it, without needing to feel 'pulled' by, or toward, it. Indeed, there's often too little time between the onset of the perceptual seeming and the formation of the belief in which to experience such a transitional sensation. For example, as S walks around having visual experiences as-of things in their environment being a certain way, they normally just passively form a corresponding set of beliefs that things are that way. This is not a consciously mediated process, but an automatic, unconscious one. However, if S does fail to believe a perceptual seeming – if, say, they have a visual perceptual seeming, but stop themselves from forming a corresponding belief, because they think themselves the victim of an optical illusion, for instance – this can cause them to feel tempted to do so, and sometimes that temptation is hard to resist. For example, even though I know that the lines in *Fig 1* are actually the same length, when I look at them I can feel consciously drawn toward believing that they are, as they seem, different. This temptation was especially strong when only recently acquainted with the illusion, where I had to continuously reaffirm to myself that, in this instance, things aren't how they appear, to avoid surrendering to my perceptual seeming. Also, while disbelieving a perceptual seeming is a particularly sure-fire way to experience this temptation, note that S can likewise experience it when they do believe said seeming, but it's at least contested by something else, such as a weak partial belief to the contrary. Here, the mere fact that S is even considering disbelieving that things are the way they perceptually seem to be can cause a reactive temptation in the opposite direction, back toward their seeming: one can feel encouraged, or pulled, toward continuing to believe that things are as they seem. Generally, then, we might say that perceptual seemings resist disbelief.

Likewise, there are several cases where contextualists' intuitions consciously tempt their subjects to believe them. Later, in §4.6, wherein I discuss the relationship between such intuitions and doxastic temptations – which, to reiterate, are distinct from doxastic dispositions – I'll list multiple instances of contextualists describing finding *p* intuitive as feeling tempted to believe it. I'll argue that, while their intuitions aren't reducible to such temptations, they are feelings that can ground (cause, or explain) such phenomena. This is, therefore, another resemblance between their intuitions and perceptual seemings. Indeed, contextualists' intuitions give rise to temptations to believe their content both when their subjects outright believe them, and when their subjects are somewhat drawn in the opposite direction, say by a rival intuition, as is the case with the Harman-Vogel paradox. In short, contextualists' intuitions resist disbelief in just the same way that perceptual seemings do.

So, the above are ten respects in which contextualists' intuitions resemble perceptual seemings in general. Following a process of elimination, and attending to these numerous similarities, when contextualists' describe their intuitions as things 'seeming' a certain way to them, by far the most plausible interpretation of their meaning is that their intuitions are perceptual seemings.

In the next section, I'll argue that contextualists' intuitions are, more specifically, perceptual seemings with an intellectual, rather than experiential, aetiology.

4.2 Intuitions as Intellectual Perceptual Seemings

In the preceding section, I identified multiple *types* of seeming, including hedged beliefs, subjective probabilities, and, crucially, perceptual seemings (or phenomenological appearances). However, there are also multiple *aetiologies* of seeming, that is, there are different sources from which seemings in general might arise. Specifically, seemings can be aetiologically experiential[46], memorial, introspective, or intellectual (c.f. Brogaard, 2013a). In other words, things can 'seem' a

[46] It's *prima facie* plausible that experiences just are (a variety of) perceptual seemings: that for it to experientially appear to S that *p* is just for it to phenomenologically appear to S that *p*, sensorily. It what follows, I'll assume this to be the case. Note, however, that Brogaard (2013a) distinguishes perceptual seemings from experiences. For her, it perceptually seeming to S that the lines in *Fig 1* are different lengths, for instance, isn't the same as S' conscious experience of those lines. Brogaard argues this on the grounds that, for example, seemings have propositional content – it seems to S that the proposition 'the bottom line is longer' is true – while experiences have concrete content (in this case, the lines themselves). Thus, for Brogaard, perceptual seemings are only grounded in, not identical with, experiences: it perceptually seems to S that *p* because S experiences the lines a certain way, for instance. This view is also compatible with other types of seeming (like hedged beliefs) being based in experiences, and, therefore, with other types of seeming having an experiential aetiology. In my account, by contrast, I'll simply deny that experiential seemings, specifically, can have said propositional content. Nothing I say in this thesis hinges on which of these two views is correct, however, since, as I'll shortly argue, the seemings that are relevant to my own investigation aren't aetiologically experiential.

certain way to S when they experience something, try to recall something, introspect on their own mental states, or just think about (or reflect upon) something, respectively[47]. The phenomenological impression, or perceptual seeming, that the lines in *Fig 1* are different lengths, for instance, has its source in sensory experience, specifically visual experience. By contrast, if S vaguely recalls that it was raining yesterday, and grounds a hedged belief or subjective probability in that fact, then it 'seeming' to S that it was raining yesterday has its source in recollection[48]. If, instead, S internally reflects upon their own mental states, and it seems to them, in whatever sense, that they believe *p*, then said seeming has its source in introspection. Finally, if S just sits and thinks about *p*, and it seems to them that *p* is (or is not) true, then said seeming is aetiologically intellectual.

I've already argued that contextualists' seeming talk (and, by extension, their synonymous intuition-talk) refers to seemings of the *perceptual* variety. In this section, I'll make the case that it refers, more specifically, to *intellectual* perceptual seemings. I'll argue for this using both positive and negative strategies. That is, I'll argue both that the relevant seemings don't have their source in experience, memory, or introspection, and that they do have their source in reflection.

First then, I'll eliminate the rival possibilities. It's clear that contextualists' seemings don't have their source in experience. Here, it's important to clarify that I'm using 'experience' to refer only to what we might call 'sensory' experiences, meaning experiences derived from the five paradigmatic bodily senses of sight, smell, hearing, touch, and taste. In the relevant literature, contextualists never describe how anything – let alone intuitive propositions – phenomenologically seems to them when they look at, smell, hear, touch, or taste it, so their seeming-talk isn't referring to experiential seemings specifically. Indeed, many of the things that 'seem' to them just aren't the sorts of things that *can* plausibly seem so sensorily. For example, contextualists often intuit propositions – such as 'ignorance', 'closure', and 'scepticism' – to be true or false and, again, what they intuit is the same as what seems to them. However, it's at least highly controversial to claim that sensory seemings can have propositional content, since it's strange to think that we can literally see, hear, smell, taste, or touch propositions, let alone sensorily experience their 'truth'. As Shafer (2013), for example, says, "the fact that some

[47] Brogaard (2013c, p.9) understands intellectual seemings as those based on what she calls 'armchair reflection', which, she says, involves bringing to bear "both a priori principles and past experience". For now, I prefer to just say that such seemings are the result of reflection on *p* very broadly construed, where that can be as minimal as *p* just flashing before S' mind for a moment and, in that instant, seeming true (or false).

[48] A seeming can also be based in more than one source, such that the different aetiological types aren't mutually exclusive. S might, for example, think it subjectively probable that *p* on the combined basis of both something they experience and something they remember. This seeming would be, aetiologically, both experiential and memorial.

proposition is true is by no means a component of the content of an ordinary perceptual experience" (p.270). Conversely, what does seem a certain way to us, experientially, are objects and states of affairs. For example, when it experientially seems to me that there's an apple on the table, what seems to me – what phenomenologically appears to me to be a certain way – is an apple and a table. What my phenomenology itself doesn't convey is a proposition – 'there is an apple on the table' – looking true.

Moreover, the propositional contents of contextualists' seemings are sometimes abstract (recall, propositions that specify no particulars). For example, general epistemic principles, such as the sensitivity principle, or CPK, are sometimes intuitive, or 'seem true', to them, and, as previously mentioned, such principles are abstract universals. By contrast, even if we conceded that experiential seemings can have propositional content, they even more plausibly can't have abstract propositional content, because what seems to us experientially is always particular *concreta*, such as tables and apples, not *abstracta*. Put simply, there has to be something concrete (or, in the case of hallucinations, for instance, something seemingly concrete) for us to experience. Thus, abstract propositions, which specify no such concrete particulars, can't be objects of experience. When *abstracta* seem to be a certain way – say, true or false – it isn't because they look, sound, taste, or smell that way, except, at most, metaphorically (c.f. Brogaard, 2012).

While these alleged differences between the content of contextualists' seemings and that of experiential seemings are controversial, and I'm not strongly committing myself to either of them, they represent further reasons to doubt that contextualists' are using seeming-talk, and hence intuition-talk, to refer to experiential seemings specifically.

It's also most unlikely that contextualists' seemings are meant to be memorial. They never once describe themselves as seeming to *recall* that 'closure' is true, for example. Rather, they simply think about that claim, and it seems true to them, then and there. Likewise, contextualists' seemings clearly aren't based in introspection, as aetiologically introspective seemings are invariably how S' *own* mental states and experiences seem when they internally reflect upon them, and a lot of the time what contextualists intuit, or what seems true to them, aren't facts about their own minds, but about, again, general epistemic principles for example.

So, in summation, contextualists' intuitions aren't perceptual seemings with an experiential, memorial, or introspective aetiology. I'll now present my positive case for the sole remaining possibility: that such seemings are, instead, of the intellectual variety.

First, we should get clear about the phenomenon in question. Traditionally, the philosophical focus has been on the perceptual use of 'seems' as it's used in reference to how things seem experientially, such as visually. However, far from all phenomenology is sensory. States like emotion also have a phenomenal aspect. To be angry or sad, for example, feels (or 'seems', in this sense) a certain way. And, most importantly, there can also be 'cognitive phenomenology', meaning how it feels to think, reflect, and the like. Intellectual perceptual seemings belong to the latter category. Sometimes, just thinking about, or reflecting upon something is sufficient for it to *phenomenologically seem* a certain way – to 'feel' true (or false), for example. When that happens, it's an intellectual perceptual seeming. So, what unifies our perceptual seemings, as a category, is that they're how things 'feel' to us, or our phenomenology, but one thing that sub-divides them is whether that phenomenology arises from experiences – the five paradigmatic bodily senses – or from somewhere else. One such alternative source is reflection, and perceptual seemings with that particular aetiology are *intellectual* perceptual seemings.

Why, then, think that contextualists' use seeming-talk – and thus intuition-talk – to report intellectual perceptual seemings, specifically? Well, while contextualists don't describe things seeming a certain way when they look at, hear, taste, smell, or touch them, talk about what they seem to recall, or report how their own mental states seem introspectively, they often do, as we've seen, describe how things seem when they reflect upon them. When they reflect upon 'scepticism' in isolation, for example, it feels false to them, but, when they reflect upon it in combination with 'ignorance' and 'closure', it seems both true and false. Moreover, the above concerns that I raised concerning the content of contextualists' seemings don't apply here. Since nothing prevents S from reflecting upon propositions, including abstract propositions, nothing prevents propositions, including abstract propositions, from feeling true or false to S when reflected upon.

Furthermore, because intellectual seemings occur to one on the basis of mere reflection, to confirm whether something intellectually seems to S the same way that it does to a contextualist, S need only reflect upon it. By contrast, to know whether something experientially seems the same way, S must have it as an object before their senses. Likewise, to know whether something memorially seems the same way to S as it does to a contextualist, S must have also encountered it in the past, so that they can compare their recollection. Finally, to know whether something introspectively seems the same way to S as it does to a contextualist, it must be a shared mental state: S must also have it 'in their heads', so that they can compare how it seems 'subjectively', or 'internally'. Therefore, a reliable test for whether a contextualist statement like 'it seems that *p* is true' is reporting an intellectual seeming, specifically, is to see whether S can determine whether *p* also seems true to them without needing anything more than to have the proposition in

question before their mind. And, indeed, contextualists don't qualify their seeming-claims by making them conditional on, say, having encountered the relevant propositions in the past, or having them 'inside' of us, as the content of our own mental states. Rather, giving the reader the propositions is meant to be sufficient for them to share the relevant intuitions. Merely giving us AI, for example, is meant to be sufficient for us to identify the seemings that constitute the Harman-Vogel paradox. Thus, contextualists behave as though the relevant perceptual seemings are, again, of the intellectual variety.

In conclusion, contextualists use seeming-talk to report *intellectual* perceptual seemings, specifically. Since they use this terminology interchangeably with intuition-talk, this is strong evidence that contextualists' intuitions are *intellectual perceptual seemings*.

In the next section, I'll argue against existing accounts of philosophical intuitions as intellectual seemings. More accurately, I'll demonstrate that such accounts fail to describe the nature of contextualists' intuitions, specifically, because they're too conservative. In place of these, I'll advocate a more liberal thesis.

4.3 Against Extant Intellectual Seemings Accounts

It's hardly unprecedented to equate (some or all) philosophical intuitions with intellectual seemings. However, in this section, I'll argue that extant intellectual seemings accounts fail to capture the phenomenon at work in contextualist arguments specifically. The common flaw with these rival positions is that they each introduce additional preconditions for an intellectual seeming being an intuition that contextualists' intuitions don't satisfy. In other words, existing intellectual seemings theories define intuitions too narrowly, or conservatively, and thus undergenerate. I'll demonstrate this by simply outlining the additional preconditions that have been posited, and arguing, in turn, why each isn't a necessary condition for being an intuition, as the contextualists, specifically, are using the term.

First, some intellectual seemings accounts claim that, for it intellectually seeming to S that *p* to qualify as an intuition, *p* must seem to S to be, not only true (or false), but *necessarily* true (or false). Brogaard (2013a), for example, claims that it intellectually seeming to S that *p* (or ~*p*) is only an intuition if "either *p* is necessary or ∼*p* is necessary" (p.212). Bealer (1998) has a similar view, however, for him, the appearance of necessity doesn't distinguish intuitions *simpliciter* from intellectual seemings more generally, but specifically philosophical intuitions from physical intuitions. Call this whole family of theories – according to which intuiting *p* (or ~*p*) means it

intellectually seeming that *p* (or ∼*p*) is *necessarily* true (or false) – the 'modal theory of intuition'[49].

There are at least three different ways that one might formulate the modal theory of intuition. On the one hand, one might claim that, for an intellectual seeming to qualify as an intuition, the proposition that seems true must actually be necessarily true or necessarily false, that is, the seeming must have necessarily true content or necessarily false content. This version of the theory appears in Brogaard, for whom a condition of intuiting *p* is, again, 'either *p* is necessary or ∼*p* is necessary'. This precondition, however, clearly doesn't apply to contextualists' intuitions, specifically. Many of the propositions that contextualists call intuitive aren't necessarily true or false, but contingently so (cf. Ludwig, 2010, p.434). The sceptical claim that S doesn't know that they have hands, for example, is meant to be intuitively false, when considered in isolation, but certainly isn't necessarily true or necessarily false, since S' situation can clearly change in ways that would add to or subtract from their knowledge. It's possible for S to know that they have hands, only to subsequently lose their hands in an accident, for example, and therefore cease to know that they have them (S can't know that they have hands if they don't have hands). More generally, we can say that many contextualist intuitions are about what S knows, and knowledge is contingent, not necessary.

On the other hand, one might claim that the content of the seeming need only be intuited as – need only 'intellectually seem' – necessarily true or necessarily false, whether it actually is or not and, here, there's a further distinction to draw between two, superficially similar views. First, one might say that the propositional content of the seeming must, itself, contain a modal operator, in other words, that the intuited proposition, *p*, must be a proposition about the necessary truth or falsity of something, and therefore have content of the form '□*q*' or '□~*q*'. The intuited proposition might, for example, be 'necessarily, all bachelors are unmarried'. However, I can't locate even a single example of an intuition with this sort of modal content in contextualist literature. Contextualists intuit that, for instance, S doesn't know that they're not a BIV, or that they have hands, not that these things are necessarily the case or necessarily not the case. Indeed, they never describe their intuited propositions using modal language. One might object that the modality is implicit, not explicit: that contextualists tacitly intuit such propositions to be necessarily true. However, given that there's no textual evidence of them doing this, such a retort would be *ad hoc*[50].

[49] See also, for example, BonJour (1998) and Sosa (2007c).

[50] This isn't to assume that contextualists couldn't have intuitions with necessarily true or necessarily false content, or that they couldn't intuit that things are necessarily true or necessarily false. I only claim that

Second, rather then claiming that, to qualify as an intuition, the content of an intellectual seeming must contain a modal operator, one might instead claim that the phenomenology of the seeming must perceptually present its content (even when that content is non-modal) as necessarily true (or false). According to this view, the propositional content of intellectual seemings, when they're intuitions, *feels*, not just true or false, but necessarily so, even if the content itself doesn't attribute necessity. However, again, there's no textual evidence of this when it comes to contextualists' intuitions, specifically. They talk of certain things feeling true or false, but never describe them as feeling necessarily so.

Moreover, in DeRose, at least, we can find direct evidence against this phenomenological claim. In his constructive account of how S comes to know (by ordinary epistemic standards) that they're not a BIV, DeRose attributes the cause and justification of S' belief that they're not a BIV to it seeming to S that they're not one, and this seeming is, again, what he elsewhere calls an intuition. He also claims that S' resultant knowledge that they're not a BIV is *a priori*, rather than based on experience. In a subsequent discussion, he addresses the fact that acquiring such *a priori*, or 'immediate' knowledge, as he sometimes calls it (2017, p.227), is often associated with it phenomenologically seeming to you that something is necessary, citing Plantinga (1993):

> "[I]t is often held that there is a characteristic phenomenology to the way we come to have *a priori* knowledge of necessary truths… And it is widely held that this phenomenology involves coming to think or its seeming to us that the truth being apprehended is *necessarily* true" (p.105).

However, DeRose observes that the kind of phenomenology he has in mind – the kind of seeming or intuition that allegedly grounds S' *a priori* knowledge that they're not a BIV – need not be the feeling that S being a BIV is *necessarily* false, but only that it's *obviously* false (absurd, or ridiculous) (2017, p.227)[51]. So, in at least some cases, contextualists' intuitions aren't something seeming necessarily true or necessarily false, but just outrageous, for example.

In summary, no version of the modal theory of intuition holds true for contextualists' intuitions, specifically.

Moving on, some intellectual seemings accounts claim that, for an intellectual seeming to qualify as an intuition, it must be 'internally correctable'. In other words, although seemings are fallible,

---

contextualists' intuitions don't *necessarily* have necessarily true or necessarily false content, or content that attributes necessity, *contra* the modal theory of intuition.

[51] DeRose compares this to the feeling described in Reid (1785, p.462).

when those that are intuitions go astray, those mistakes can always, in principle, be detected, and corrected, merely through reflection. For example, BonJour (1998) says:

> "Further reflection on the very state or process that led to the mistaken result is capable of revealing that it was a mistake and of replacing it with the correct result" (p.116).

However, contextualists' intuitions don't always (or, indeed, often) display internal correctability, as I've already shown. To reiterate, if faults in these intuitions could not only be identified but also corrected merely through further reflection, then it would be possible for the contextualists to make their intuitions cohere with what they take to be true: to amend their perceived intuitive errors. Moreover, we would expect contextualists to reconcile their intuitions with their beliefs if they could, if only to relieve the cognitive dissonance. As we've observed, however, many contextualists' intuitions blatantly contradict what they take to be true (c.f. Huemer, 2007, p.38). Therefore, they don't display internal correctability[52]. So, being internally correctible isn't a precondition of being an 'intuition', as contextualists are using the term.

Another view states that only an intellectual seeming that arises 'immediately' – rather than following, or as a result of, lengthy, conscious reflection upon its content – qualifies as an intuition. Brogaard (2013a), for instance, says:

> "Not all intellectual seemings are intuitions… 'It intellectually seems that p' expresses an intuition only if the mental state it refers to forms immediately upon considering p and not as a result of extensive, explicit reasoning" (p.9)[53].

However, while some contextualist intuitions do, perhaps, arise immediately – the intuition that 'scepticism' is false, when it's considered in isolation, is a plausible example, since that proposition immediately 'strikes us' as false when we read it, without the need for any explicit, conscious reflection on the matter – some do not. For instance, the contrary intuition that 'scepticism' is true most often and most readily emerges (according to the contextualists' own claims) when we emulate the sceptic's reasoning, that is, when we hold the premises of AI in our head, in our conscious 'working memory'[54], and infer, or at least imagine inferring, 'scepticism'. It's that process of consciously simulating AI that gives us the intuition that 'scepticism' is true,

---

[52] It's still technically possible that contextualists' intuitions would correct themselves if only their subjects would put even more effort into the task, but that's a very *ad hoc* way of saving the internal correctability thesis. As it stands, we simply have no reason to think that additional effort would make any difference.

[53] BonJour (1998) expresses a very different view. For him, the fact that an intellectual seeming follows careful, explicit consideration of its propositional content doesn't exclude it from being an intuition. Indeed, such consideration is, for him, necessary for said intuition to justify beliefs.

[54] I'll say more about this in the next chapter (in §5.2), where I discuss the burgeoning view that philosophical intuitions are automatic, or 'type 1' judgements.

and gives rise to the cognitive dissonance of the Harman-Vogel paradox. That procedure is certainly explicit, since we *consciously* reason through the argument, and protracted, since it takes more a flash for the intuition to emerge – one (especially the uninitiated) must spend at least a few moments performing the multi-premise inference. Therefore, the sceptical intuition, that 'scepticism' is true, unlike the anti-sceptical one, that 'scepticism' is false, precisely isn't the immediate, unreflective impression we have of that proposition, but something that comes about only after explicit, protracted reflection on the matter. So, contextualists' intuition-talk doesn't only refer to immediate intellectual seemings.

Yet another view is that an intellectual seeming with content *p* only qualifies as an intuition if it bestows *prima facie* justification on the belief that *p* (Brogaard, 2013a; Huemer, 2007, 2008; Pryor, 2005). Initially, it doesn't seem too implausible that this might be a precondition for being an 'intuition' in the sense that contextualists use the term, specifically. After all, DeRose, for example, we may recall, advocates 'phenomenal conservatism', and claims that intuiting *p* – or it seeming to them that *p* – grants S defeasible justification to believe *p* (and both he and Cohen frequently take intuiting, or it seeming that, *p* to justify philosophical theses, as we've seen). However, neither DeRose nor Cohen ever says anything to the effect that intuitions justify beliefs *by definition*, such that, if they did not so justify, then they wouldn't be intuitions at all. I suspect that, if one were to convince Cohen or DeRose that what they call 'intuitions' don't possess such justificatory power, it wouldn't stop them from using the word 'intuitions' to describe them, but merely change their views about their epistemology. In other words, contextualists would still call their intuitions 'intuitions' even if they didn't justify belief, so bestowing *prima facie* justification isn't a precondition for being an 'intuition', as they're using the term.

Finally, some claim that it intellectually seeming to S that *p* only qualifies as an intuition if it intellectually seems to S that *p* solely on the basis of S' understanding of *p* (Boghossian, 2009; Sosa, 2007c, p.101[55]). However, if it were true that contextualists' intuitions, specifically, always arose purely on the basis of their understanding of their content (or 'conceptual competence'), then their intuitions would never depend on bringing to bear anything besides the intuited proposition itself: reflection upon *p*, and its conceptual content, would always be sufficient to have the intuition that *p* is true (or false). As we've seen though, this is far from the case. Again, the intuition that 'scepticism' is true, at least, doesn't arise purely from consideration upon

[55] As I'll discuss later, Sosa thinks that intellectual seemings are just doxastic temptations. So, for him, to say that intuitions are intellectual seemings with a basis in conceptual competence is to say that they're doxastic temptations with a basis in conceptual competence.

'scepticism', but upon, at minimum, the premises of AI as well, and so that intuition certainly isn't based *solely* on one's understanding of 'scepticism'.

In summation, extant intellectual seemings accounts of intuition posit various criteria for distinguishing intuitions from intellectual seemings more generally. However, contextualists' intuition-talk doesn't exclusively refer to something that satisfies even one such precondition. Consequently, prior intellectual seemings accounts of philosophical intuitions fail to accurately characterise contextualists' intuitions, specifically, because they're all too restrictive, or conservative, about which intellectual seemings intuition-talk refers to. Indeed, as we can see above, most intellectual seemings theories of intuition fail on multiple accounts, because they propose more than one of the discussed criteria. To circumvent the same mistake, my own view will be less restrictive: I'll instead assume that contextualist intuitions are intellectual (perceptual) seemings *simpliciter*[56].

In the next section, I'll explain how my positive account of contextualists' intuitions can be reconciled with, and even help us to make sense of, the many, inconsistent ways that they describe said intuitions in the relevant literature.

4.4 Reconciling Disparate Textual Evidence

I promised earlier that my constructive account of contextualists' intuitions would respect the evident connection or similarity between said intuitions and the various other phenomena with which they are sometimes conflated in the literature, namely, outright beliefs, credences, doxastic dispositions, and ordinary linguistic practices (or facts about them). In this section, I argue that a rational reconstruction of contextualists' intuition-talk as referring to intellectual perceptual seemings does exactly that. For, although my interpretation of contextualists' intuitions is superficially at odds with some of the many other ways in which they choose to describe them, my reading does help us to understand the connection between intuitions and these various 'nearby' phenomena and, in turn, why contextualists chose to describe their intuitions, misleadingly, as such things, specifically. Indeed, it actually goes a long way to explaining the confusing, incoherent presentation of the intuitive data in contextualists' arguments.

[56] There may, of course, be other respects in which contextualists' intuitions can be delineated from intellectual (perceptual) seemings in general. I only mean that I reject all existing attempts to identify which proper sub-set of such seemings are intuitions, and that, for now, I'll assume a minimally restrictive account, for simplicity.

To elaborate, while intellectual perceptual seemings are neither partial nor outright beliefs, since they don't entail either, they do still, normally, ground, or explain, such doxastic states. When it perceptually seems to S that *p*, this will typically cause S to believe *p*, or at least make them more confident that *p*. When, for instance, it perceptually seems to S that there's an apple on the table, they'll normally believe, on that very basis, that there's an apple on the table. Consequently, our perceptual seemings are bound up quite tightly with our beliefs and credences. This, in turn, creates a strong connection between intuitions and doxastic dispositions, since, to reiterate, the degree of S' disposition to believe *p* is just a measure of how often they believe *p*, so, if perceptual seemings ground, and thus correlate with, beliefs, then they also ground, and thus correlate with, dispositions to believe. The same basic idea also applies to speech-talk and attribution-talk. If this is referring to what people believe, then the very same connection holds here, between what they 'say' or 'attribute' and how things 'seem' to them. If, instead, it's literally referring to ordinary language, then we can reiterate that what people sincerely utter will track what they believe, and thus, again, how the world perceptually seems to them.

The correlation between intuitions and these various other phenomena explains why competent English speakers casually vacillate between, say, describing intuitions as their beliefs (or 'judgements') and confidences (what they find 'plausible', 'compelling', 'persuasive', and so on), on the one hand, and how things 'seem' to them, on the other. Put simply, the casual interplay of different terms – like 'belief', 'judgement', 'confidence', how things 'seem', and what we 'say' – perhaps motivated by the desire to circumvent repetitive prose, works for most intents and purposes, because of the strong correlation between those various states. Usually, one does, for instance, believe what perceptually seems to them to be the case, so fluctuating between the two states, while it might confuse us somewhat, is normally innocent and harmless. This is probably why contextualists, describing philosophical intuitions in what is, evidently, a fairly casual, unreflective way, will employ all of this terminology.

However, when we encounter intuitive paradoxes, such as the Harman-Vogel paradox, this normally harmless interplay of different terms can become more confusing. Since we have conflicting intuitions about, say, AI, it's unlikely that we'll believe, be confident in, assert, or be disposed to believe, all of them. Thus, in such atypical cases, intuitions – or intellectual perceptual seemings – start to come apart from the various other states with which they can normally be interchanged, and the standard manner of describing such intuitions in these many different ways can, therefore, lead to contradictions. Here, one might not, for example, believe what perceptually seems, or 'feels' to them to be the case, but still find it natural to describe such an intuition, unreflectively, as a 'judgement'. Hence, it's not surprising that contextualists use misleading, though totally normal terminology when describing Harman-Vogel intuitions, nor that,

in such abnormal instances, such language creates confusion over what the relevant 'intuitive data' actually are.

4.5 Perceptual Seemings as Presentations

In section §4.1, I argued that contextualists' intuitions are perceptual seemings, and that such seemings are presentational, to wit, a kind of mental state that presents the world, to their subject, as being a certain way. There, I specifically argued that said intuitions are presentational in the way that Bengson (2015) uses the term. Bengson also advocates a view (the 'quasi-perceptual' view) according to which philosophical intuitions are presentational. However, he construes his presentational thesis as a rival to the seemings thesis (*Ibid.* p.729). More precisely, he claims that intuitions are presentations, but that presentations and seemings are different kinds of state, and that, therefore, intuitions aren't seemings. Though he doesn't address contextualists' intuitions specifically[57], if Bengson is correct in thinking that, in general, seemings are distinct from presentations, then my own view, according to which said intuitions are both, is mistaken.

It this section, I defend my claim that contextualists' intuitions are presentational states more thoroughly, as I promised earlier. Specifically, I'll address Bengson's concerns about reconciling a presentational theory of intuition with a seemings account of intuition[58]. I'll argue that such a synthesis can, after all, be achieved. Combining these accounts produces a unique, hybrid theory of philosophical intuition.

Bengson has argued against conflating presentations and seemings on two grounds. First, he claims that advocates of the seemings theory of intuition tend to treat the relevant 'seemings' as primitive, and offer little to no account of what such states actually are. By contrast, Bengson has provided a positive characterisation of presentations. As such, we shouldn't see his claim, that intuitions are presentational, as an addendum to an existing view, but as a superior, more informative alternative.

I concede that past advocates of the seemings theory of intuition have often treated the relevant seemings as primitive, or at least failed to provide much in the way of a positive metaphysical

[57] Also, to reiterate, Bengson's own project isn't to discern the general meaning of intuition-talk in analytic philosophy, but merely to identify a genuine, presentational kind mental state, which some philosophers, including him, often call 'intuition' (*Ibid.*, p.734). Nonetheless, his claim that presentations and seemings are distinct is still incompatible with my own account of contextualists' intuitions, specifically.

[58] In fact, Bengson admits to being 'friendly', in principle, to the possibility of framing his account of presentations as an elaboration upon the metaphysics of seemings (*Ibid.*, p.730). He just has some reservations about such a reconciliation, which I'll address herein.

characterisation of them. However, that's partly why I devoted so much space in this chapter to elaborating upon the properties of said seemings. In §4.1, I identified no fewer than ten properties of these states, many of which help us to delineate the relevant species of seeming from other, nearby phenomena – such as hedged outright beliefs and subjective probabilities. So, the account I've offered is hardly primitive or uninformative. Moreover, perceptual seemings in general display many of the very properties that Bengson himself uses to identify presentational states – such as being involuntary, gradable, and tempting belief in their content, for example – strongly suggesting that said seemings are indeed presentational in the sense that he has in mind.

Further still, I think that, if anything, Bengson's observation gets things backwards: it's the presentational account, not the seemings account, which fails to posit a referent for the word 'intuition'. To elaborate, perceptual seemings are an actual phenomenon, but there's no such thing as presentation*s tout court*. Rather, some states have the property of being presentation*al* and, clearly, not all of them are intuitions. Some, for instance, are sensory experiences. So, Bengson's quasi-perceptual model of intuition still leaves us wondering: in what does the property of presentationality adhere in this instance, if not in sensory experiences? In other words, just what *are* intuitions? By contrast, the seemings account at least tells us what intuitions are – seemings – not just one of their properties. Attaching the property of 'presentational' to perceptual seemings is just another step to explicating their metaphysical features. It tells us that perceptual seemings – such as intuitions – are presentational, in addition to the other characteristics of theirs that we've identified.

Bengson's second argument against conflating presentations with seemings is more complex. He claims that there's at least one, crucial difference between presentational states and seemings. Specifically, seemings[59] are always 'explicit'. By this, he means that the content of one's occurrent seeming is always consciously available to, and fully articulable by, its subject: "one is able at the time to formulate explicitly — out loud or in one's head — the way things seem [to them]" (*Ibid.*, p.730). Notably, this point isn't argued for, but taken for granted. By contrast, he says, one might not be in a position to articulate fully or precisely how things are presented to them.

The first reply that might occur is that nothing about Bengson's two claims – that presentations aren't always explicit, and seemings are – is incompatible with the presentationality of seemings. To elaborate, even if we concede that not all presentational states are explicit, this doesn't entail

[59] Bengson doesn't specify which sense of 'seeming' he has in mind when he makes this claim. Indeed, he criticises the seemings account to philosophical intuitions without disambiguating the kind of seeming state that's relevant to that discussion. As such, it's not obvious how he understands such views, and to which state, exactly, he's attaching the property of being explicit.

that none of them are. After all, there are multiple types of presentational state, even according to Bengson (intuitions and sensory experiences, to name but two). Thus, there might be a certain sub-class of presentations that *are* explicit, and those might be (the relevant kind of) seemings. That is, while not all presentations are explicit, like seemings are, all seemings (of the relevant kind) might nonetheless be explicit presentations.

However, what's at stake here is whether intuitions, specifically, are explicit. If they're not, and Bengson's two claims are true, then, while they might be presentations, they can't be seemings. This inference is valid even if there are some explicit presentations, and indeed even if that category includes seemings. So, charitably, we should assume that Bengson's intended claim is, more specifically, that not all *intuitions* are explicit. More precisely, Bengson's position is that, while one might not always be able to articulate, fully, how things are presented to them by intuition, they will always be able to articulate, fully, how things seem to them and, therefore, intuitions aren't seemings. To demonstrate, Bengson invites us to imagine cases of the following form:

> A logic professor and an undergraduate student are both looking at one of De Morgan's laws, and both finding it intuitively true. In this case, what is being *presented* to them by intuition is exactly the same: the equivalence rule. However, only the professor is in a position to fully articulate, out loud or in their head, the exact logical equivalence rule that they're intuiting. The student, on the other hand, cannot; at best, they can articulate only a very vague idea of what the rule says. The upshot is that the professor and the student have exactly the same intuition – what is presented to them as true is the same – but how things *seem* to them must differ. Specifically, it seems to the professor that something very specific is true, while, to the student, it only seems that a vague *something* is true (cf. *Ibid.*, p.730-731)[60].

I have two objections to Bengson's view that intuitive presentations aren't explicit, while seemings are. First, I challenge the unargued claim that seemings are explicit. For S to be able to *fully* articulate how things seem to them, they must *know*, precisely and completely, how things seem to them. More precisely, S must know exactly what the content of their seeming is, not that such content is true. Therefore, for seemings to *always* be explicit, they must be 'luminous' in some sense. That is, it must always be possible for a subject to know what the content of their

[60] The example that Bengson discusses in most detail is 'Kleinblue'. In this case, an artist familiar with a certain shade of blue (Kleinblue) and an amateur who is not are both having it *presented* to them, by their visual experience, that there is something Kleinblue on the canvas, but it only *seems* to the artist that it's, specifically, Kleinblue, while, to the amateur, it only seems that it's *very dark blue*. Since 'Kleinblue' deals with a case of visual presentation, not intuition, it isn't as useful for a discussion of intuitions.

seemings are, it cannot be possible for S to be wrong about how things seem to them (or, at least, such mistakes can't be intractable, and should be easily rectified), and it can never be the case that S knowing how things seem to them is, in principle, impossible. However, there's growing doubt about the possibility of such luminous states in general (see, especially, Williamson (2000) for the most influential version of the classic anti-luminosity argument, and Srinivasan (2015) for a significantly updated version). The burgeoning view is that subjects can fail to know, or be wrong about, the contents of any mental state they have and, moreover, that, for any such state, it's possible for the subject to be unable to know what its content is, even in principle. Given that the worry about luminosity isn't specific to any particular kind of mental state, but a concern about the logical (or, in the case of Srinivasan, empirical-psychological) possibility of luminous states in general, there's no reason that these concerns wouldn't apply with equal force to the view that seemings, in particular, are explicit. To clarify, while we don't have to assume that the anti-luminists are right, it's become exceedingly difficult to convincingly argue for (let alone casually assume) the luminosity of any particular state. Accordingly, we shouldn't suppose that the difference between presentations and seemings lies in one of them being luminous[61].

My second objection to Bengson concerns the kind of examples that he uses to demonstrate his view. In the De Morgan's case above, for instance, we're invited to agree that the logic professor's intuition and the student's intuition – what's being presented to them – are the same, while what seems to them differs. I can concede the latter claim – what seems true to the professor is, for example, something much more specific and vivid than what seems true to the student – but not that their intuitions are the same. I would say that, in such cases, the student and the professor simply have different intuitions. For one, the content of their intuitions seems dissimilar. The former has the impression that just *what they're looking at* is true, the latter the impression that *a very specific equivalence* holds. I can discern no, non-arbitrary reason to think that, here, the content of the two intuitions is exactly the same, and the professor just has a clearer or sharper grasp of what that content is. Indeed, Bengson says that presentations are gradable, and can have more or less vivid or exact content, in the same way that experiences can. Therefore, it's entirely consistent with his own characterisation of presentational states that the professor's intuitive presentation of the equivalency rule is just more vivid than the student's. In other words, rather than thinking that the professor and the student have identical presentations, or intuitions, with

[61] Bengson might reply to this first objection that, even if seemings aren't strictly speaking luminous, there's still an obvious enough, though somewhat weaker, epistemic difference between seemings and presentations that he's trying to pick up on. I don't see yet how the specifics of such a reply might be developed.

the same content, we can say, using Bengson's *own* criteria for describing presentational states, that there are two, different intuitions here, with differently precise content.

But there's one more, crucial step to my counterargument. Even if how things are presented to the professor and the student, and how things seem to them, both come apart, this doesn't show that their seemings and presentations are the same states. However, we can note that the two phenomena differ in exactly the same respects. Just as what seems true to the professor is something more exact than what seems true to the student, what's intuitively presented to the professor is now likewise something more exact than what's intuitively presented to the student. Hence, the two phenomena – how things seem and what's intuited – align perfectly. From here, nothing prevents us from simply equating (the relevant kind of) seemings with (certain) presentations. To clarify, whether we want to say that that presentations, like seemings, are explicit (that both subjects are able to fully articulate how things seem to them, and how things seem to them is just how they're presented as being), or that neither presentations nor seemings are explicit, we have every reason to think that the two phenomena are identical. Either way, then, Bengson's claim about seemings being explicit doesn't drive a wedge between the seemings account of intuitions and the presentational account.

Therefore, *contra* Bengson, we can reconcile the presentational and seemings accounts of philosophical intuitions, as I have done.

4.6 The Phenomenology of Intellectual Perceptual Seemings

In §4.1, I argued that contextualists' intuitions, at least, are perceptual seemings. Perceptual seemings are, to reiterate, how things seem phenomenologically, or, in other words, how things literally feel. In §4.2, I argued that the relevant perceptual seemings have their source in reflection: they're *intellectual* perceptual seemings. However, at this stage, all we know is that some intuitions are 'feelings', broadly construed, grounded in reflection. It has yet to be explicated what kind of feeling these intuitions constitute. In this section, I expand upon my constructive account by investigating the phenomenology of intellectual perceptual seemings (hereafter, just 'intellectual seemings') more closely.

One possibility is that intellectual seemings are, more precisely, doxastic temptations, as I defined them in Chapter 3. According to this thesis, experiencing an intellectual seeming is just feeling attracted to assent to (or deny) a proposition. This view has been put forward by, most notably,

Sosa (2007a, Ch. 3)[62], Boghossian (2009, p.116), and Williamson (2007). The latter, for example, equates an intellectual seeming with a 'conscious' 'temptation' to believe, saying:

> "I am aware of no intellectual seeming beyond my conscious inclination to believe Naïve Comprehension, which I resist because I know better. I can feel such an inclination even if it is quite stably overridden, and I am not in the least danger of giving way to temptation" (p.217).

If it's true that intellectual seemings are just doxastic temptations, then it follows that at least some philosophical intuitions, such as those reported by contextualists, are likewise such. In fact, the view that philosophical intuitions are doxastic temptations is hardly unprecedented. Van Inwagen (1997), for example, identifies intuitions as things that:

> "make certain beliefs attractive to us, that 'move' us in the direction of accepting certain propositions without taking us all the way to acceptance" (p.309).

Here, intuitions don't, in themselves, take us 'all the way to acceptance', but merely urge, or 'move' us toward belief.

Indeed, there's some textual evidence to suggest that contextualists' intellectual seemings, or intuitions, specifically, might be doxastic temptations. For one, both DeRose and Cohen frequently describe their intuitions in terms of 'pushes' and 'pulls'. When DeRose, for example, claims that his EC can explain the intuition that, even when epistemic standards are high, 'scepticism' is false, he sometimes presents his thesis as being able to explain, more specifically, the 'pull' that the denial of 'scepticism' exerts upon us, as he does here:

> "[T]he very strong pull that [~scepticism] continues to exert on (at least most of) us even when the standards are high is explained in the manner outlined…" (2017, p.31, my emphasis).

Indeed, when discussing 'ignorance', DeRose goes as far as describe his intuitions about it as *conscious* 'pushes' and 'pulls' toward *believing or disbelieving* it, specifically:

> "Though I could always *feel* a strong *intuitive pull toward accepting* this premise [ignorance], I found it far from compelling, and I also *felt* a significant opposing *intuitive push toward denying it*" (*ibid.*, p.55, my emphasis).

[62] This, of course, represents the alternative interpretation of Sosa's dispositionalism that I alluded to in Chapter 3. According to this reading, Sosa equates intuitions, not with doxastic dispositions, but with conscious 'attractions to assent' (see, for instance, Sosa (2007a, p.54)).

Cohen (1998), meanwhile, describes the intuition that 'scepticism' is true in the same terms, namely, as a conscious lure toward accepting it:

> "[W]hen in the grips of a sceptical argument, we feel the strong intuitive pull of saying that our knowledge ascriptions are mistaken" (p.302, my emphasis).

So, both our contextualists describe their intuitions, or intellectual seemings, in terms of doxastic temptations: as feeling tempted to believe or disbelieve propositions.

However, in this section, I'll argue that this textual evidence is misleading, and that the reported intuitions (or intellectual seemings) aren't conscious temptations to believe (or disbelieve) propositions, but feelings that said propositions are true (or false). To elaborate, the two phenomena at least have different *objects*: the latter is directed at the proposition *p*, or is how that proposition, itself, feels, while the former is directed toward the mental action of judging it, to wit, it's that act, itself, feeling attractive, in some sense (cf. Kornblith, 1998). I hope to clarify this distinction with an example. Freshly cooked chips feel hot to the touch, and I feel tempted to eat them. However, what feels hot, and what feels tempting, are two, quite different things. While what feels hot are the chips themselves, what feels tempting is the act of placing them in my mouth and masticating. The temptation is directed toward a complex action that the chips, at best, form only a part of. Likewise, while the feeling that *p* is true (or false) is simply and entirely about *p* itself, the temptation to judge that *p* (or ~*p*) is directed at a more complex mental action, of which *p*, at best, forms only a part.

p feeling true (or false) can, and often will, ground (cause, or explain) a temptation to believe (or disbelieve) it. Therefore, we can explain the interplay of intuitions and doxastic temptations in contextualist literature in much the same way as I earlier explained their proclivity to describe said intuitions as outright beliefs, credences, doxastic dispositions, and ordinary speech acts, or facts about them. That said, this section marks the final transition from exegesis to theory. I will, going forward, assume that my interpretation of contextualists' intuition-talk, as referring to intellectual seemings, is accurate, on the basis of the reasons already provided, and merely seek to elaborate upon that phenomenon. That is, I'll provide a theory of intellectual seemings – as feelings of truth, not doxastic temptations – rather than an interpretation of what contextualists' take such things to be.

I have three arguments supporting my position that intellectual seemings are feelings that propositions are true, not temptations to believe them, which I call 'the argument from overgeneration', 'the argument from analogy with experience', and 'the argument from truth-aptness'. In the remainder of this section, I'll present each argument in turn.

First, 'the argument from overgeneration'. Put simply, it intellectually seeming to S that *p* shouldn't be conflated with S experiencing the doxastic temptation to believe *p*, because it's possible for S to feel tempted to believe *p* without it seeming to S that *p*. Suppose that, for example, S hopes that *p* is true. As I explained earlier, hopes aren't representational, let alone presentational states. Accordingly, in merely hoping that *p*, it's neither represented in S' mind, nor presented to S' mind, that *p* is true. Thus, in merely hoping that *p*, it doesn't 'seem' to S, in any sense, that *p*. To clarify, a precondition of it seeming to S that *p*, in any sense, is *the world seeming* a certain way to S, and thus the world being at least represented in, if not presented to, S' mind as being such and such a way. For example, while what delineates perceptual seemings from other kinds of seeming is that they're the world *phenomenologically* seeming a certain way to S, this doesn't entail that all phenomenal occurrences are perceptual seemings, because there can be phenomenal occurrences, or feelings, that aren't representational[63]. Consequently, S hoping that *p*, because it's a nonrepresentational state, isn't sufficient for it seeming to S that *p*. Yet, hoping that *p* can, in itself, give rise to doxastic temptations. The desire-like state of hoping that *p* is true can result in, or perhaps even partly consist of, a desire-like urge to believe that it is. Indeed, in the phenomenon commonly called 'wishful thinking', S is biased toward certain cognitive endorsements because they're drawn toward believing what they hope to be true (cf. Bengson, 2015, p.726). So, S can feel inclined to believe *p* without it seeming to them that *p*. Therefore, equating intellectual seemings with doxastic temptations 'overgenerates', that is, it posits intellectual seeming where there are none.

Second, 'the argument from analogy with experience'. Over the years of writing this thesis, I've become so well acquainted with the Müller-Lyer illusion that, when I look at *Fig 1*, I no longer experience any conscious temptation to believe that the bottom line really is longer than the top one. I did, of course, reliably feel such a temptation before I indoctrinated myself against it, but now I don't have to exert any cognitive effort whatsoever to stop myself from believing that the world is the way that it perceptually seems to be in this instance. Nonetheless, the lines still perceptually seem to me to be of different lengths: that's still how they seem phenomenologically, or how they, quite literally, look. Thus, in this case, it perceptually seems to me that *p*, even though I'm not at all tempted to believe *p*. Indeed, even if one suspects that I simply must experience *some* degree of temptation to believe my perceptual seeming as-of the lines being different lengths, it's certainly not as though, as I become less tempted by said seeming, it perceptually *seems less* to me that one line is longer. It seems, in this sense, just as

[63] Indeed, I'll imminently argue that doxastic temptations themselves are, plausibly, nonrepresentational phenomenal events, that is, they don't represent their objects as being any particular way, and thus can't be right or wrong about them.

strongly and vividly that the lines are of different lengths as it did when I first looked at them. So, either way, degrees of doxastic temptation don't align with degrees of perceptual seeming. Thus, the measure of it perceptually seeming to S that *p* isn't just the measure of S' temptation to believe *p*. There's no reason to think that intellectual perceptual seemings, specifically, would behave any differently (that they wouldn't be analogous to visual experiences). In other words, there's no reason to suspect that *p* would intellectually seem any more or less true to S, in the perceptual sense, the more or less tempted S feels to believe it, or *vice versa*. Therefore, the measure of it intellectually, perceptually seeming to S that *p* likewise isn't just the measure of S' temptation to believe *p*. We might alternatively call this 'the argument from undergeneration': just as there can be cases where S feels tempted to believe *p* without it seeming to them that *p*, there can be cases where S doesn't feel tempted to believe *p*, where it does seem to them that *p*, in the relevant sense, and, therefore, equating intellectual seemings with doxastic temptations sometimes fails to posit intellectual seemings where they exist, or 'undergenerates'.

Third**,** and finally, the argument from truth-aptness. Intellectual seemings are 'truth-apt', meaning, their content is either true or false. As we've seen, DeRose takes the intellectual seeming, or 'intuition', that knowledge requires sensitive belief, for example, to be false. He could be wrong in his assessment, of course, but, either way, it's clear enough that the said seeming is *either* true or false: knowledge either requires sensitive belief, or it doesn't. Owing to the aforementioned ambiguity between mental states and their content, what we strictly mean by this is that the *content* of intellectual seemings is either true or false; the content of DeRose's intellectual seeming, for example, is a proposition like 'sensitive belief is a precondition of knowledge', and that proposition is either true or false, depending on what knowledge requires[64]. By contrast, temptations don't have truth-apt content. To elaborate, rather than propositions, the contents of temptations are *actions*, or, what S is tempted to do. These can be physical actions, such as eating chips, or, as is the case with doxastic temptations, mental actions, such as judging. Actions, however, are neither true nor false. It doesn't make sense to say that the act of eating

[64] Again, this isn't to assume that all perceptual seemings have propositional content. Intellectual perceptual seemings, however, do, or at least can.

chips, for instance, is 'true' or 'false'[65]. Therefore, intellectual seemings aren't doxastic temptations, or indeed temptations of any variety (cf. *ibid.*, p727)[66,67].

Now, one might respond that temptations are desires, or at least somewhat desire-like, and thus involve something like 'the guise of the good'. To elaborate, desires arguably represent things as being good, valuable, or otherwise desire-worthy, or otherwise represent the act of acquiring or using them as good, valuable, pleasurable, or otherwise worth doing (see, for instance, Tenenbaum (2013) for a discussion of this notion). My desire to eat chips, for example, might represent chips as, say, delicious, or otherwise worthy of being eaten, or, alternatively, represent the act of eating them as, say, pleasurable. One interpretation of that alleged phenomenon is to say that desires simply are a kind of belief (so-called 'deliefs'); the desire to eat chips, for instance, is the belief that chips are tasty, or that eating them would be pleasurable. Even without going that far, however, if desires somehow represent things as being a certain way, then they might, indeed, have truth-apt propositional content[68]. The desire to eat chips, for example, might have, as its content, a proposition like 'chips are delicious' or 'eating chips feels good', and that content would be either true or false, depending on whether chips really are delicious, or whether eating them really does feel good. Therefore, if temptations are desires, or similar to desires in this regard, then temptations could have propositional, and thus truth-apt content. So, intellectual seemings could still be doxastic temptations.

However, even if I did concede that doxastic temptations have truth-apt content, they would still differ from intellectual seemings in at least one, crucial respect. To elaborate, on this reading, the temptation to believe *p* would have, as its content, a proposition that ascribes to *p* properties that would make it good to believe, or otherwise belief-worthy, or otherwise attribute positive characteristics to the act of judging *p*. Truth might be one such property: if a proposition is true, that makes it more belief-worthy than a false one would be, and a judgement with a true outcome is better than one with a false outcome. However, there are many other properties that

---

[65] At best, an action is either actualised (or instantiated) or not, can be 'right' or 'wrong', strictly in the sense that it's moral or immoral (or otherwise normatively correct or incorrect), and can be authentic/honest or inauthentic/dishonest.

[66] To clarify, under a doxastic temptations account of intellectual seemings, the contents of said seemings are the judgings themselves, not the propositions that S is tempted to judge. Thus, judging itself would have to be truth-apt and, as an action, it isn't. The fact that the propositions that S is tempted to judge are truth-apt is incidental.

[67] More generally, said seemings have representational content (of which propositional content is but one variety) – they represent their object as being this or that way, and can be right or wrong in that assessment – while temptations are merely attractions toward actions, and plausibly don't represent those actions as being any particular way.

[68] Again, this doesn't necessarily follow, because representations, in the broad sense that I define them in this thesis, need not have propositional content. However, if desires are representational, then they're at least potentially propositional.

can contribute to a proposition being belief-worthy, or to a judgement being good to perform. The proposition in question might be supported by S' evidential *corpus*, or at least consistent with it, for example. A temptation might assign these features to *p* without explicitly representing *p* as true. Moreover, this is so far assuming that S is perfectly rational, and isn't tempted to believe *p* on obviously bad grounds, such as the fact that, if *p* were true, that would be good for S, or because other people, maybe emotionally close to S, believe it too. Consequently, the truth or falsity of *p* – the proposition that S is tempted to believe – won't always be sufficient (and, sometimes, may not even be necessary) to render the propositional content of the temptation itself true or false, respectively: sometimes, *p* would also have to possess, or lack, certain other features. Therefore, it's possible for *p* to be true, and the temptation to believe it (or rather, the content of such a temptation) be false, and for *p* to be false, but the temptation to believe it (or rather, the content of such a temptation) be true. By contrast, it's almost platitudinous that, if sensitive belief isn't a precondition of knowledge, for example, then the intellectual seeming that the sensitivity principle is true, for instance, is false. Nothing more or less is required to render the content of said seeming true or false. This suggests that, in difference to the doxastic temptation to believe *p* (or ~*p*), the content of the intellectual seeming that *p* (or ~*p*) is always and only that *p* is true (or false). That's at least the simplest explanation for why the truth-value of *p* is always necessary and sufficient to fix the truth-value of the seeming. Put simply, the content of doxastic temptations – even supposing, controversially, that it's propositional, or truth-apt – would have to be more varied, and sometimes more complicated, than the content of intellectual seemings. So, again, intellectual seemings differ from doxastic temptations.

So, on the basis of these three arguments, we shouldn't conflate contextualists' intuitions, or intellectual (perceptual) seemings, with doxastic temptations. However, to reiterate, perceptual seemings can, and often do, ground such temptations (I'm sometimes tempted to believe that the lines in *Fig 1* are of different lengths *because* they perceptually seem that way, for example), and that helps to explain the occasional tendency, such as we see in contextualist literature, to conflate talk of intuitions with talk of temptations or 'pulls'.

Instead of the temptation to believe (or disbelieve) *p*, then, I propose to understand the intellectual (perceptual) seeming that *p* (or ~*p*) as the feeling that *p* is *true* (or false). I'll say a lot more about what that constructive account amounts to in the next chapter.

4.7 Chapter 4 Conclusion

In conclusion, according to the most charitable reconstruction of their ambiguous terminology, contextualists' intuition-talk refers to intellectual perceptual seemings, which are presentational states, and not reducible to doxastic temptations.

In the next chapter, I'll expand further upon my constructive account. Having discerned that at least some philosophical intuitions are being cited in an evidential capacity, in some of the most influential arguments in modern epistemology, and that said intuitions are intellectual perceptual seemings, I'll argue that such seemings are a sub-category of what some psychologists call 'epistemic feelings', or 'e-feelings'. I'll ultimately conclude, therefore, that such e-feelings have been instrumental to shaping modern epistemology. In Chapter 6, I'll briefly explore some of the possible implications of this discovery.

Chapter 5 Intuitions as Epistemic Feelings

In this thesis, I've thus far argued that contextualists, at least, use intuition-talk to cite evidence for their theories, and that contextualists, at least, use such terminology, roughly and perhaps unknowingly, to refer to intellectual perceptual seemings, specifically. The argument for EC that such seemings are used in, namely 'the argument from anti-scepticism', has contributed significantly to the popularity of that thesis. Therefore, intellectual perceptual seemings have been instrumental, via the burgeoning influence of EC, for example, in shaping modern epistemology[69]. In this chapter, I argue that intellectual perceptual seemings aren't *sui generis*, as many philosophers are wont to assume, but a sub-category of what some psychologists call 'epistemic feelings' (hereafter, 'e-feelings'), and that, consequently, said feelings have had, and are by all accounts continuing to have, a profound impact on our discipline.

In §5.1, I outline the general characteristics and functions of e-feelings. Then, in §5.2, I make the case for equating intuitions – or intellectual perceptual seemings – with specific e-feelings, namely, certain feelings of rightness (or FORs). In §5.3, I anticipate a possible objection to my claim: e-feelings, it's commonly supposed, are metacognitive, and metacognition, it's sometimes assumed, requires metarepresentation, while intuitions have first-order, or non-metarepresentational content. I respond that, contrarily, e-feelings aren't metarepresentational. In the next chapter, I'll consider some of the profound ramifications of my discovery that e-feelings lay at the methodological heart of some of the most influential arguments in modern epistemology.

5.1 Epistemic Feelings

In this section, I outline the general characteristics of e-feelings, their function, and significance.

5.1.1 General Characteristics of E-Feelings

Broadly construed, epistemic feelings (or e-feelings) – otherwise known as 'noetic feelings' (Arango-Muñoz, 2014) or 'epistemic emotions' (Carruthers, 2017) – are feelings about epistemic matters. This includes, for example, feelings about the accuracy of a proposition, about the likelihood of something being true, or about the validity of an argument. To help clarify, two

[69] Going forward, I'll sometimes simply call intellectual perceptual seemings 'intuitions' (and *vice versa*), though bear in mind that I still only meant to refer to some philosophical intuitions, such as those cited in the arguments for EC. Again, we have good reason not to overgeneralise in our metaphilosophical claims.

illustrative and oft-discussed examples of e-feelings are the 'feeling of rightness' (the FOR) (Gangemi *et al.*, 2015; Prowse Turner and Thompson, 2009; Thompson *et al.*, 2011), also known as the 'feeling of certainty', or the 'feeling of confidence' (Arango-Muñoz, 2014), and its corresponding, negative sensation, the 'feeling of error' (the FOE) (Gangemi *et al.*, 2015), also known as the 'feeling of wrongness' (Thompson and Morsanyi, 2012), or the 'feeling of uncertainty' (Arango-Muñoz, 2014; Dunlosky and Metcalfe, 2009; Carruthers, 2017)[70]. A subject, S, might, for instance, feel that a proposition, *p*, is true, or that it's false, or feel that an argument for *p* is valid, or that it's invalid: the (epistemically) positive sensations are feelings of rightness (or FORs), the negative ones feelings of error (or FOEs). When S hears the sentence 'it will rain tomorrow', for example, it might feel true to them (or probable), or feel false (or unlikely). Likewise, S might feel that 'the argument from anti-scepticism' for EC, for instance, is valid, or that it's invalid.

For the most part, research into e-feelings has focused on their role in metamemory and metaperception, that is, on how these experiences might be triggered by, and help to guide, recollection and perception, respectively (Reder and Ritter, 1992; Schunn *et al.*, 1997; Vernon and Usher, 2003; Koriat, 1993, 1995; Koriat *et al.*, 2004; Whittlesea, 1993; Whittlesea & Williams, 2001). Indeed, until the last decade, it was almost unprecedented to consider how e-feelings might relate to thinking and reasoning more generally, beyond their role in regulating memory (Thompson *et al.*, 2011). Lately, however, it's been recognized that FORs and FOEs are actually ubiquitous, and can be triggered by S merely 'thinking' something, in either the strong or weak sense of the word. To elaborate, on the one hand, S might actually believe, infer, or otherwise conclude that *p* ('think' *p*, in the strong sense), and feel that their judgement is, indeed, correct, or that it's mistaken, or that their reasoning toward *p* was valid or invalid, sound or unsound. For example, when solving problems that give rise to cognitive illusions, subjects may detect bias or other sources of unreliability in their cognitive faculties, get the feeling that something, though they don't always know exactly what, has gone wrong in their reasoning, and experience a feeling of error, a FOE, toward their own judgement (Piattelli-Palmarini, 1994; Thompson and Johnson, 2014). On the other hand, however, S might merely read or hear *p*, or reflect upon it only hypothetically, or have it just 'cross their mind' ('think' *p*, in the weak sense). Such weak thoughts are likewise sufficient to trigger e-feelings toward their contents. In other words, S doesn't have to actually believe *p* in order to feel that *p* is true (or false) when they 'think' it: it can also feel true (or false) when they just read or hear it, for instance. Likewise, S doesn't have to actually

[70] Some other examples of e-feelings include the feeling of surprise (Dunlosky & Metcalfe, 2009), the feeling of curiosity (Carruthers, 2017), and the feeling of interest (*ibid*.).

perform a piece of reasoning to feel that it's valid or invalid, sound or unsound, because they can merely read, simulate, imagine, or reflect upon that reasoning, say by encountering it in a philosophy paper, and have it feel a certain way to them when they do.

E-feelings can also be about actions, rather than about arguments or propositions. In such cases, FORs and FOEs, for example, are feelings of certainty and uncertainty, respectively, about performance, or, more precisely, about the probability of an action succeeding. For instance, when imagining themselves trying to jump a wide gap, S might feel that they would be unlikely to succeed, and thereby experience a FOE toward said action (Carruthers, 2017, p.70). E-feelings can be about both physical actions, such as jumping a gap, and mental ones, such as judging, or retrieving an item of information from memory. For instance, another commonly discussed example of an e-feeling is the 'feeling of knowing' (or FOK) (Arango-Muñoz, 2014). The FOK is a FOR directed, specifically, at the act of recollection, prior to any information having actually been recalled. To elaborate, when asked a question, S may feel that they'll be able to answer it, even before they've retrieved the answer from memory. That is, they may consciously sense that they know the answer, even before they consciously known what the answer is, and thus feel that the act of memory retrieval is likely to succeed:

> "Consider, for example, the experience that you underwent in school when the teacher asked a question to one of your classmates and you felt that you would have been able to answer it, even before you or your classmate could retrieve the correct answer" (*Ibid.*, p.195)[71].

Typically, S experiences a FOK because they can very quickly bring certain details of the answer to mind, such as the first letter, or recall closely related memories in the vicinity of the answer, and this makes them confident in their ability to quickly narrow-down the sought after fact (Dunlosky and Metcalfe, 2009). Conversely, when S feels they cannot retrieve the information, they experience a 'feeling of not knowing' (Glucksberg and McCloskey, 1981), which is a FOE directed, specifically, at the act of memory retrieval. That is, they feel that said action is unlikely to succeed. These particular e-feelings are, then, feelings of certainty and uncertainty directed at mental actions themselves, rather than at the propositional contents of mental states.

[71] The term 'feeling of knowing', or FOK, is used slightly differently in some places, for example, in Gangemi *et al.* (2015), Efklides (2006), Koriat (2000), and Nelson and Narens (1980). There, the term refers to when people fail to recall an item of information, but nonetheless feel that they would successfully do so on a later test. The difference here is that, instead of the FOK being the feeling that an imminent recall will be successful (as in classroom style situations), it's the feeling that one does know the information, and so will be able to recall it on future occasions, even if they failed to do so in this instance.

The above examples also help to demonstrate how e-feelings can arise at different stages of cognition (Arango-Muñoz, 2014). Some, such as the FOK, appear at the predictive stage, or before an action: S feels that they will be able to retrieve the desired information, that is, that a future mental action would or will be successful. Some other e-feelings, however, occur at the intermediate stage, or during an action (Koriat, 2000; Thompson, 2009; Walsh and Anderson, 2009). For example, the FOK is closely related to another, particularly distinctive e-feeling, the 'tip-of-the-tongue' feeling (or TOT) (Brown and McNeill, 1966; Arango-Muñoz, 2014; Dunlosky & Metcalfe, 2009). The TOT is the feeling that S has when they experience a FOK, but then fail to quickly retrieve the sought after information. Typically, S does retrieve at least some details of the answer, such as its first letter[72], giving them the feeling that the whole answer is 'on the tip of their tongue', and will arrive imminently, motivating them to continue the recall effort, rather than give up. Here, the TOT e-feeling is directed toward an occurrent cognitive process: it's the feeling that the current retrieval attempt will succeed, if S simply perseveres. Finally, there are e-feelings that arise at the post-evaluative stage, or after an action (Thompson, 2009; Koriat, 1993, 1995, 2000). This includes, for example, the aforementioned feeling that S' own judgement was somehow mistaken, or that something went astray in their reasoning process. Here, S experiences a FOE toward their own beliefs after forming them, or toward their own reasoning after preforming it.

5.1.2 The Functions and Significance of E-Feelings

E-feelings, and the systems that produce them, are often assumed to be metacognitive. Gangemi *et al.* (2015), for example, casually describe e-feelings as 'metacognitive experiences' (cf. Thompson *et al.*, 2011, p.108). Now, as Nagel (2014) observes, "researchers have been able to produce strikingly divergent theories of the domain they label 'metacognitive'" (p.710). However, the general idea is that metacognition constitutes one cognitive subsystem monitoring and regulating another in the same subject (Dunlosky and Metcalfe, 2009; Flavell, 1979; Nelson and Narens, 1990). In other words, metacognition is when one mental system monitors, detects, or is otherwise sensitive to, features of another mental system, or of its outputs, evaluates said mental phenomenon on the basis of such features, and, if necessary, takes steps to correct, (re)calibrate, or otherwise guide and control it. Proust (2014), for example, defines metacognition thusly:

> "Metacognition is the set of capacities through which an operating cognitive subsystem is evaluated or represented by another subsystem (p.4)".

[72] This only makes sense, given that a FOK is usually triggered by S already having some details of the answer, or some closely related piece of information, in mind.

Thus, we can understand metacognitive mechanisms as monitoring and control subsystems, and metacognitive mental states as the outputs of said mechanisms, being assessments of other subsystems. Indeed, the latter is exactly how many researchers describe e-feelings. They are, according to Thompson *et al.* (2011), for example, akin to what Koriat *et al.* (2006) would call a "subjective assessment of one's own cognitive processes" (p.38). That is, e-feelings are evaluations of S' own mental states and activities, arising from metacognitive monitoring and control subsystems.

Now, claiming that e-feelings are metacognitive, though commonplace, isn't entirely uncontroversial, and I'll say a lot more about metacognition, and its relationship to e-feelings, in §5.3. For now, however, we can roughly characterise the received wisdom thusly: an e-feeling is essentially an epistemic assessment of some mental process, its output, or the propositional content of its output, based on higher-order facts, rather than first-order considerations. To clarify, here, 'first-order' considerations means world-oriented facts, while 'higher-order' considerations means subject-oriented facts, to wit, facts about S' own mental states and processes. Correspondingly, 'first-order' cognition means thinking based on first-order facts, while 'higher-order' cognition means thinking based on higher-order facts.

For example, suppose that S infers *p* from some first-order premise, *q*. They might then experience a further, positive feeling toward *p* (a FOR), not on the basis of *q*, but on the higher-order grounds that the process of inferring *p* from *q* was particularly fluent (fast and easy) (Koriat, 2000), that *p* itself was easy for S to parse (understand) (Whittlesea, 1993; Whittlesea & Williams, 2001), or the frequency with which *p* has previously been encountered by S (Reder, 1996). S might, for instance, conclude that 'scepticism' is true, on the grounds of 'ignorance' and 'closure', and then experience a further, positive feeling toward 'scepticism' (a FOR), not on the fist-order basis of 'ignorance' and 'closure', but on the higher-order grounds that inferring 'scepticism' from these premises was particularly fast and easy, or fluent. Likewise, S might feel that they'll be successful in recalling the name of their previous landlord (that is, experience a FOK) on the higher-order grounds that they were able to retrieve some details of the desired information very quickly, such as the first letter of the name. In such cases, S' e-feelings are assessments of their own mental states and processes, or of their contents, based on facts about said states and processes. Such metacognitive assessments, performed by monitoring and control mechanisms, can then serve to guide, correct, or motivate S' first-order cognition. A FOR might, for example, reinforce S' belief in 'scepticism', while a FOK might motivate them to recall the name of their former landlord from memory, rather than resorting to alternative information-gathering strategies.

Indeed, whether we regard them to be specifically metacognitive or not, e-feelings certainly aren't epiphenomenal; rather, they're known to play a significant role in explaining both bodily (de Sousa, 1987; Döring, 2003) and mental (Peacocke, 2007, 2009; Proust, 2009a) action. As Arango-Muñoz (2014) summarises:

> "Philosophers and psychologists have become interested in these experiences not only because of their puzzling character, but also because of the role they might play within the cognitive architecture and mental life of a subject" (p.193).

For example, when asked a question, how long, and how carefully, a subject rethinks their initial answer, before giving their final verdict, depends on, among other things, how strong their FOR toward that initial answer was (Thompson and Johnson, 2014). Studies also suggest that, in game show style situations, wherein participants are tasked to answer questions before their rivals, subjects rely on their FOKs to sense whether they'll be able to answer successfully. In other words, if they experience a strong enough FOK, contestants will press the buzzer to answer the question before they've actually recalled said answer from memory (Reder, 1996). FOEs, on the other hand, can motivate a subject to avoid or abandon a particular action, including a cognitive action, either before or during the attempt (this obviously only refers to cases wherein the feeling of uncertainty is process-directed and either predictive or intermediate, not answer-directed and post-evaluative). If S feels that they likely won't, for example, be able to recall a piece of information from memory, of if they're currently experiencing difficulty doing so, then they might consider alternative means of finding the answer they seek, as the current strategy feels like a poor investment of time and energy (Carruthers, 2017, p.70-74).

Of course, as Nagel (2014, p.713) keenly observes, when it comes to determining the functions of e-feelings, it's crucial to isolate the effects of the affective cues themselves from those of the mental states, contents, and processes that give rise to them. After all, the fact that certain mental processes and outputs coincide with e-feelings, and precede changes in S' thinking or behaviour, doesn't show, in itself, that the e-feelings themselves caused said change, rather than the mental phenomena that they accompanied. However, there's growing empirical evidence that it is, indeed, e-feelings themselves that influence our thoughts and behaviour. For example, it's known to be fairly easy for people to think of six situations in which they behaved 'very assertively and at ease', and very difficult to think of twelve. If you ask one group of experimental participants, group A, to come up with 6 such examples, and group B to come up with 12, then group B will subsequently rate themselves as less assertive than group A, despite having just (successfully) thought up twice as many examples of cases wherein they acted assertively. Schwarz *et al.* (1991) conclude from this that the added difficulty of thinking up 12 cases directly

affected the self-evaluation of group B: if it's hard for you to think of examples of assertiveness, then you must not be very assertive, or so they assumed. Thus, the feeling of difficulty (or disfluency) influenced people's self-evaluations. Other empirical work suggests that feelings of ease and difficulty, or fluency and disfluency, give rise to FORs and FOEs, respectively, which, in turn, directly affect judgement (see, for example, Alter and Oppenheimer, 2009).

Likewise, there's evidence to the effect that how intensely subjects experience a FOR toward their initial answer is what directly contributes to the duration and degree of further reflection. For instance, further reflection is more likely to be engaged (or engaged to a greater extent) when conflict (the activation of processing pathways that interfere with each other, such as when assimilating contradictory information, see Koriat and Levy-Sadot, (1999)) is detected in the reasoning task that produced the initial answer (Bonner and Newell, 2010; De Neys and Glumic, 2008; Thompson *et al.*, 2011; Thompson and Johnson, 2014). However, the evidence shows that it's not conflict *per se* that encourages further reflection, but the lower FOR (or higher FOE) that accompanies such conflict. In a series of studies by Thompson and Johnson (2014), it was found that "markers of analytic engagement… were more pronounced for conflict than non-conflict items" but "those individuals whose FOR judgments were not sensitive to conflict did not show this pattern" (p.238). In other words, the extent of further analytic engagement with a question is determined by the presence and strength of a FOR, not by the presence of conflict: without experiencing a low FOR, subjects don't engage in further reflection, even in cases of conflict. Again, this makes the e-feeling itself directly efficacious.

So, e-feelings are crucial to explaining how people behave and think. Despite this, they're rarely incorporated into theories of mind and action, especially in philosophy (de Sousa, 2008; Efklides, 2006; Schwartz and Metcalfe, 2010). This thesis will do something to correct this trend, specifically by integrating e-feelings into a positive account of intuition, and, by extension, into theories of philosophical method.

5.1.3 E-Feelings and Emotions

As I previously mentioned, e-feelings are sometimes called 'epistemic emotions'. This is because they're often thought to share characteristics with more paradigmatic emotions. For example, according to classic 'intentionalist' or 'representationalist' theories, emotions are, specifically, feelings *about* an object or state of affairs (one that may or may not actually exist) (Crane, 2001; de Sousa, 1987; Dretske, 1995; Goldie, 2000; Tye, 1995). In other words, emotions have content: they're feelings about, or toward, *something*. This is one respect in which emotions are thought to differ from feelings more generally. If S is sitting in a slightly uncomfortable chair, for instance,

their feeling of discomfort isn't 'about' anything: it's merely a phenomenal feature of their existence, or how it presently feels to be S. By contrast, when S experiences the emotion of anger, they're angry at, about, or toward something, or someone. Put simply, emotions, but not feelings in general, are contentful. E-feelings are likewise contentful[73]. FORs, for instance, are feelings about propositions (that they're true, or probable) or actions (that they're likely to succeed). Therefore, in this respect, at least, e-feelings do resemble emotions.

However, on some views, emotions, specifically, are dualistic. On the one hand, an emotion involves 'feelings towards', being the aforementioned feelings directed at the content, or object, of the emotion, and, on the other, it involves 'bodily feelings', being feelings in, or of, S' own body (Goldie, 2000, 2002). For example, the emotion of hunger involves S feeling the contractions of their own stomach (bodily feelings) and feeling desire, temptation, or positivity toward food (feelings towards) (Arango-Muñoz, 2014). However, while e-feelings certainly involve feelings towards – feelings about propositions or actions – it's less obvious that they all involve bodily feelings. There are no particular bodily sensations associated with the FOR and the FOE, for example. S might, for instance, feel that De Morgan's Laws are true, but not experience any palpable physiological reaction when having this thought. We do, of course, sometimes describe such a FOR as a 'gut feeling', but it's unclear how literally we should interpret this language: introspectively, at least, there's no particular gastrointestinal phenomenology associated with De Morgan's Laws feeling true. Moreover, while some e-feelings do activate motor plans, such as facial expressions (Prinz, 2004; Carruthers, 2017) (surprise, for instance, is typically accompanied by the raising of one's eyebrows), it's not obvious that we consciously experience these as 'bodily feelings'. It's not clear, for instance, that S feels the raising of their eyebrows when they're surprised in the same sense that they feel intestinal discomfort when they're hungry. So, e-feelings, unlike emotions, aren't dualistic.

Emotions are also thought to contain both evaluative and descriptive elements: they're "structured by a synthetic unity of 'factual' and 'evaluative' attribution'' (Roberts, 2009, p. 218). In other words, emotions report how the world is (the 'factual' attribution), but also how 'good' or 'bad' that thing, or state of affairs, is (the 'evaluative' attribution). The emotion of hunger, for example, conveys both the factual information that S' stomach is empty, that they require food, and so on, alongside a negative evaluation: this condition is bad, and should be relieved. This is to be contrasted against a non-emotional experience, such as a plain perceptual experience as-of there being an apple on the table. Said experience might present it to S that there is such an

[73] Note, therefore, that my account of intuitions, according to which they're identical with certain e-feelings, is compatible with intuitions having content.

apple, without attributing any positive or negative status to that object or state of affairs; it might simply convey the 'factual' content that the apple is there, without any 'evaluative' component (Seager, 2002; Tye, 2008).

Initially, it might seem that e-feelings are likewise both descriptive and evaluative, carrying factual information about what their object is, as well as an epistemic assessment of it. A FOR, for instance, conveys a specific proposition, and that said proposition is true, or likely, or a specific action, and that said action is likely to succeed (Proust, 2012; Carruthers, 2017). However, the evaluative component of e-feelings is quite different from the evaluative component of emotions.

Call the evaluative component of an emotion its 'valence', being the feeling of positivity or negativity directed toward the object of said emotion. The positive and negative valences involved in most emotions are best likened to 'affects'. Affects are feelings of pleasantness or unpleasantness (Carver, 2003), which can motivate behaviour, such as approach or avoidance (Carver and Scheier, 1998). When S experiences hunger, for example, they feel a positive affect toward food, or toward the act of eating it (food, or eating, feels pleasant, which motives them to pursue it) and a negative affect toward their empty stomach (the emptiness feels unpleasant, motivating them to alleviate it). Now, there's certainly a positive correlation between information that we find pleasant to process and information that we feel is true. For example, processing fluent information is more likely to be subjectively pleasant for us than processing disfluent information (Zajonc and Rajecki, 1969), and fluent information is also more likely to evoke strong FORs, that is, feel true[74] (Whittlesea, 1993; Whittlesea and Williams, 2001; Koriat, 2000).

However, it's possible to teach subjects to reverse the e-feelings that they have in response to fluency and disfluency, such that they experience FORs toward disfluent information instead. For instance, information presented in a hard-to-read colour contrast, being more difficult to process, is, by default, more likely to seem false (trigger a FOE), while easier-to-read information seems comparatively more plausible (triggers a FOR). But, if experimental participants are given feedback on a cognitive task that correlates fluent answers with false ones, and *vice versa*, then they learn to associate greater fluency with falsity, reporting FORs toward disfluent information instead in follow-up tasks (Unkelbach, 2007). Nothing about this suggests that hard-to-read text becomes pleasant to read: that, we can imagine, remains a physically uncomfortable task. So, people can learn to distrust pleasant fluency, and trust unpleasant disfluency. Therefore, it's

[74] This actually applies to both *process fluency* (the fluency with which information is processed, which is affected by things like how easy to read the font it is written in is, how hard it is to hear it spoken, and so on), and *answer fluency* (the fluency of the information itself, which is affected by things like how well it fits with our existing beliefs and expectations).

possible to experience a positive e-feeling, such as a FOR, toward unpleasant information, and a negative e-feeling, such as a FOE, toward pleasant information. So, the affect of something, and the evaluation attributed to it by e-feelings, can come apart. Thus, we cannot simply conflate the evaluations attributed by e-feelings with affects, and, in this regard, they once again differ from emotions.

Moreover, it's possible to experience an e-feeling sans any positive or negative affect toward its content at all. After all, whether information triggers an e-feeling or not has little to do with whether S experiences any pleasantness or unpleasantness toward it. S might be neutral; they might feel that an item of information – say, one of De Morgan's Laws – is correct, but this fact might fail to stir any particularly pleasant or unpleasant sensations in them, if they simply don't care that it's true, for example. The evaluative component of their e-feeling cannot, therefore, be an affect, which is absent in such cases.

Indeed, by contrast to emotions, which embed affects, or feelings of pleasantness or unpleasantness, e-feelings embed evaluations of epistemic properties. That is, e-feelings convey information about the accuracy, reliability, probability, validity, or soundness, for example, of their object, rather than its pleasantness. This is one, clear respect in which e-feelings differ from paradigmatic emotions. In fact, one might even say that the epistemic properties that e-feelings attribute – truth, falsity, and so on – are factual (or descriptive), rather than evaluative. Saying that *p* is true, for example, sounds more like reporting on how things are, rather than how good or bad it is that they are that way. Thus, it might not even be appropriate to call e-feelings 'evaluative' in the first place, further undermining the alleged resemblance between e-feelings and emotions. At least, it's a stretch to say that e-feelings are evaluative in just the same sense as paradigmatic emotions are.

None of this is to assume that any or all of the aforementioned theories of emotion are correct. I only mean to observe that, under some popular views of emotion, e-feelings may not, after all, resemble them. At least, the similarity between the two phenomena isn't, on reflection, obvious or uncontroversial. For that reason, I'll choose to call the relevant phenomena epistemic 'feelings', and forego describing them in emotional terms. Furthermore, this means that nothing about my forthcoming account of intuitions – according to which they're identical with certain e-feelings – entails that intuitions are emotions. At least, any such claim would have to rely on a more controversial theory of emotions than the ones just presented.

To recapitulate, e-feelings are epistemic assessments of actions or propositions, can be triggered by S 'thinking' something, in the strong or weak sense, can arise at different stages of cognition, are strongly associated with metacognition, and arguably differ from emotions.

In the next section, I'll argue that at least some philosophical intuitions, specifically intellectual perceptual seemings, are identical with certain e-feelings, specifically, certain varieties of FORs and FOEs.

5.2 Intuitions as Epistemic Feelings

In this section, I'll argue that certain e-feelings, namely, specific FORs and FOEs, are type-identical with intellectual perceptual seemings and, therefore, with certain philosophical intuitions, such as those cited in 'the argument from anti-scepticism' for EC.

Indeed, it's hardly unprecedented to call e-feelings, especially FORs and FOEs, 'intuitions'. Thompson, *et al.* (2013), for example, state that the ease with which fast, biased responses occur to us often creates a strong 'intuition' that they're correct, referring, specifically, to the accompanying FOR that quick responses typically generate. Gangemi *et al.* (2015) likewise call the FOR an 'intuition' (p.383), and elsewhere refer to the FOE as "an intuition of error" (p.394)[75,76]. Admittedly, most contributors to the literature on e-feelings avoid using intuition-talk, at least pervasively. Indeed, Arango-Muñoz (2014), for instance, explicitly chooses not to call e-feelings 'intuitions' in his work, precisely because, in light of the burgeoning philosophy of intuition, doing so would imply a very loaded and controversial identity claim. Nonetheless, the evident naturalness of describing these e-feelings as intuitions already lends some credence to the view that certain kinds of intuition, at least, are e-feelings. And, of course, this terminological overlap comes from the other direction as well. To reiterate, it's very common for philosophers to claim that intuitions consist of a particular phenomenology, or *feeling* (Parsons, 1995; Cappelen, 2012; Thompson *et al.*, 2011; Bealer, 1998), and this feeling is specifically described as a feeling of 'rightness' or 'correctness' (Parsons, 1995; Plantinga, 1993), just like the FOR is.

So, psychology calls FORs and FOEs intuitions, and philosophers call intuitions things feeling right (or true). Despite this, no robust defence of the position that philosophical intuitions are e-feelings has hitherto been laid out and, consequently, the thesis has yet to be established as a significant contender in the sub-discipline of metaphilosophy. In this section, I'll remedy this by providing some further considerations in favour of taking such a view seriously.

---

[75] A slight complication here is that Gangemi *et al.* (2015) also, in places, use 'intuition' to refer to Type 1 processes and/or answers (in the manner that has become common in some psychology literature), rather than to the FOR that often accompanies them, which they recognise is a distinct phenomenon. This shouldn't mislead us.

[76] For further examples of intuition-talk being used in reference to e-feelings, see, for instance, Gladwell (2007) and Volz *et al.* (2008).

First, in §5.2.1, I'll argue that e-feelings are, like experiences, perceptual seemings. Then, in §5.2.2, I'll argue that some FORs and FOEs have an intellectual aetiology. I'll conclude that said FORs and FOEs are, therefore, intellectual perceptual seemings, and thus type-identical with a certain class of philosophical intuitions.

5.2.1 E-Feelings as Perceptual Seemings

There are many respects in which e-feelings resemble perceptual seemings. Indeed, as I'll now argue, such e-feelings bear the same ten resemblances to perceptual seemings as said seemings do to contextualists' intuitions. In other words, I'll argue that e-feelings are perceptual seemings on the same grounds as I argued that said seemings are the referents of contextualists' intuition-talk.

Phenomenal. E-feelings, like perceptual seemings, are phenomenal.

Perhaps most obviously, e-feelings are phenomenal. They are, after all, *feelings* that propositions are true or false, that arguments are sound or unsound, and so on (Carruthers, 2017). Indeed, this is, recall, one respect in which they're thought to resemble emotions: both e-feelings and emotions are, or are accompanied by, a "conscious, phenomenal, or qualitative experience" (Arango-Muñoz, 2014, p.194). In other words, it feels a certain way to have an e-feeling, just as it feels a certain way to be angry.

Contentful. E-feelings, like perceptual seemings, have content.

As I explained earlier, while some philosophers distinguish emotions from feelings more generally by saying that only the former have content, e-feelings, whether they're emotions or not, are certainly contentful. When S experiences a FOR, for example, *something* feels right to them, whether it's a proposition, or an instance of, actual or imagined, reasoning. Again, e-feelings are sometimes described as 'emotions', for bearing this, among other resemblances, to more paradigmatic emotional states, and the reason I opt to call them 'feelings' instead of 'emotions' is not because they lack content, but because of other, potential dissimilarities between the two phenomena.

Representational. E-feelings, like perceptual seemings, represent the world.

E-feelings can be correct or incorrect, relative to the world. The FOR, for instance, is the feeling that, for example, a proposition, *p*, is true, or likely, or that some argument is valid, or sound, and that feeling is correct or incorrect, depending on whether *p* really is true or likely, or whether the said argument really is valid or sound. In other words, the loosely 'evaluative' component of an e-

feeling can be right or wrong in its assessment of its object, and thus, like perceptual seemings, 'represents' its object as being a certain way, in the sense that I've defined representation in this thesis.

Non-Factive. E-feelings, like perceptual seemings, are fallible representations.

Like their perceptual seemings, S' e-feelings can lead them astray. Consider, for example, the 'believability bias' (Evans and Curtis-Holmes, 2005; Tsujii and Watanabe, 2010). When S assimilates an argument with a conclusion that they already believe, they're much more likely to feel that said argument is sound. However, an argument with a believed conclusion will typically feel sound even when it's obviously invalid (when it, for example, affirms the consequent). Thus, the believability bias gives rise to misleading e-feelings. In other words, e-feelings can be inaccurate representations of the world. Therefore, e-feelings are fallible, or non-factive.

Presentational. E-feelings, like perceptual seemings, are presentational.

People sometimes describe e-feelings as 'judgements'. Thompson *et al.* (2011), for example, casually refer to FORs as 'reasoning judgements', while Proust (2012) occasionally calls such phenomena 'metacognitive judgements'. However, while a few, such as Nussbaum (2001) and Solomon (1993, 2004), have previously defended a doxastic account of certain e-feelings – according to which FOKs, for instance, are type-identical with certain judgements (here, read 'beliefs') – it's much more, and increasingly common to construe e-feelings in non-doxastic terms. Many commentators, for example, prefer to describe e-feelings as 'experiences', rather than 'judgements' or 'beliefs' (Gangemi *et al.*, 2015), or otherwise choose to keep e-feeling altogether distinct from such doxastic phenomena (Arango-Muñoz, 2014; Nagel, 2014). Moreover, while there's some indication of equivalent third-person phenomena, such as the 'feeling of another person's knowing' – which is regarded as the third-person counterpart to the FOK – researchers investigating such things class them as 'judgements', not 'feelings', and, for that very reason, choose to distinguish them from e-feelings proper (Jameson *et al.*, 1993). Meanwhile, some, such as Carruthers (2017), for example, talk of metacognitive judgements relying on, rather than being, e-feelings. Indeed, Carruthers explicitly says that e-feelings might fail to ground such judgements, and thus that "epistemic emotions don't presuppose… judgments of any kind" (p.63).

In fact, even some of those aforementioned philosophers who do use doxastic language to describe e-feelings, such as Thompson *et al.* and Proust, can be said to only mislead their readers by doing so. While Thompson *et al.* (2011) sometimes refer to e-feelings as metacognitive judgements, they also refer to them, just as often, as metacognitive 'experiences' (on p.136, for instance), and, in an earlier paper, Thompson (2009) explicitly distinguishes between the

metacognitive judgement itself and the e-feeling that produces it, saying: ‘‘it is the interpretation of that feeling or affective response that produces a judgment’’ (p.181). In Proust (2012), some terminological slippage is also apparent. While she does often talk about ‘metacognitive judgements’ in the same breath as e-feelings, strictly speaking these judgements aren’t, on her view, the e-feelings themselves, but something that’s sometimes based on such feelings. Indeed, according to Proust, what delineates one category of metacognitive judgement (‘experience-based’ metacognitive judgement) is that these judgements, in particular, are based on e-feelings. Not all metacognitive judgements are based on e-feelings though, and thus not all are experienced-based. Nor, most crucially, does an e-feeling entail the presence of a metacognitive judgement, since it might fail to ground one. So, e-feelings are certainly not, on her view, metacognitive judgements themselves.

So, most theorists, on reflection, regard e-feelings as non-doxastic, and the prevalence with which some philosophers describe e-feelings as ‘judgements’ can make it seem as though a doxastic account of this phenomenon is more common than it really is.

Indeed, there are many ways in which e-feelings differ from judgements, or beliefs, and these are the same respects in which contextualists’ intuitions were shown to differ from doxastic phenomena in §3.1. Some of these differences, such as the phenomenal nature of e-feelings, we’ve already seen, while others, such as their recalcitrance, will soon become apparent. Moreover, it’s just generally understood that S can experience an e-feeling without holding any corresponding belief in its content. As Proust (2012), Nussinson and Koriat (2008), and Unkelbach (2007), for example, have all observed, if a subject thinks that one of their FORs or FOEs, for instance, has been caused by an irrational bias, such as the aforementioned believability bias, or is otherwise unreliable, they can treat it as undiagnostic of genuine accuracy or error, and curtail the influence of that feeling upon their further, ‘downstream’ judgements. In other words, S can experience an e-feeling without endorsing its content. Thus, e-feelings don’t entail, and thus cannot be identical with, beliefs in their content.

Accordingly, e-feelings, like perceptual seemings, pass the smell test for being presentational: they’re representational, as we’ve already seen, and, as I’ve just argued, they’re non-committal, meaning, they don’t entail belief in their content. In other words, they’re how the world is ‘given’ to S, not how S ‘takes’ the world to be.

Recalcitrant. E-feelings, like perceptual seemings, are recalcitrant.

When S takes themselves to have good evidence that one of their e-feelings is misleading – such as when they have reason to think that a FOR was the result of bias – and fails to acquiesce to it,

typically, the e-feeling itself remains (Nussinson and Koriat, 2008). That is, e-feelings don't go away even when they run counter to their subject's evidential *corpus* and beliefs. They are, therefore, like perceptual seemings, insensitive to evidence, or 'recalcitrant'. Also, this is at least partly because, like perceptual seemings, e-feelings cannot be changed or undone merely by reflecting upon them, or upon the evidence against them, and thus a subject cannot bring their e-feelings into agreement with their evidence, even if they're inclined to do so. For instance, if a proposition, *p*, feels right to S, but S believes that *p* is false, they usually cannot stop feeling that *p* is right, no matter how hard or sincerely they try. Proust (2012), for example, calls this the "immunity to revision of noetic feelings" (p.252). In other words, e-feelings, like perceptual seemings, aren't 'internally correctible', and this contributes to their recalcitrance.

Involuntary. e-feelings, like perceptual seemings, are involuntary.

E-feelings are an example of what psychologists call 'type 1 judgements' (Stanovich and West, 2000; Evans, 2008; Evans and Stanovich, 2013; Kahneman, 2011; Sloman, 2002; De Neys, 2012), meaning, phenomena[77] that arise from type 1 thinking[78]. To elaborate, type 1 thinking is, by definition, *automatic*, rather than deliberate (Thompson and Johnson, 2014). If, for example, S saw me hold a glass bottle high above the ground and let it go, they would, under normal conditions, automatically form the belief that it will fall to the floor and break (and perhaps react accordingly, say by trying to catch it), rather than getting to that conclusion *via* a period of controlled, conscious reflection[79]. In cases like these, S' judgement is more like a reflex – a spontaneous response to the visual stimuli of the bottle falling – than a deliberate mental action. Such automatic, or 'type 1' inferences often involve the application of 'heuristics', to wit, rules of thumb that typically, though not invariably, produce the correct answer (Kahneman, 2003). For

[77] I choose to say 'phenomena' here, since what psychologists call type 1 'judgements' aren't always judgements in the sense that epistemologists, for example, typically use that term, namely, beliefs, or acts of belief formation. E-feelings, for instance, are, as I stated earlier, generally understood to neither be, nor entail, belief in their content, and yet, e-feelings are thought to be examples of type 1 'judgements'. In other words, the relevant psychology literature, 'type 1 judgements' just means phenomena (including, at times, beliefs) that arise from type 1 thinking.

[78] Type 1 and type 2 thinking also go by many other names in the literature, such as 'system 1' and 'system 2' (Stanovich, 1999), 'heuristic' and 'analytic' (Evans, 2008), 'associative' and 'rule-based' (Sloman, 1996), 'automatic' and 'controlled' (Schiffrin and Schneider, 1977), or 'thinking fast' and 'thinking slow' (Kahneman, 2011), respectively. Each of these other names, however, implies some further assumption about the nature of the cognitive divide. For example, the terms 'heuristic' and 'analytic' imply that all type 1 thinking involves the application of heuristics, which may not be true, as type 1 may also be able to execute, at least, some fairly simple logical deductions, such as *modus ponens* (De Neys, 2012; Handley *et al.*, 2010; Pennycook *et al.*, 2013). Likewise, it's known that it can take longer to apply some automatic heuristics than it does to perform some simple, deliberate deductions, so type 1 thinking isn't necessarily faster than type 2 thinking, *contra* to what is implied by the 'thinking fast, thinking slow' nomenclature (Thompson and Johnson, 2014, p.220). I therefore elect to just use the terminology 'type 1' and 'type 2', since this language is the most neutral.

[79] This is inspired by a similar example that appears in De Cruz (2014, p.5).

instance, in this example, S is likely following some heuristic along the lines of 'if someone lets go of a glass bottle in mid-air, it will fall to the floor and break'[80]. Such rules, normally drawn from past experience, can be applied very quickly, and with little to no effort (Nagel, 2011; De Cruz, 2014; Evans and Stanovich, 2013; De Neys, 2006). Following them is, therefore, somewhat akin to taking 'mental shortcuts' that circumvent the need for slower, more effortful reflection. Finally, automatic judgements are typically based on unconscious grounds, rather than explicit ones. In the case above, for example, the heuristic that S relies upon is not something that they consciously bring to mind before applying it, nor do they necessarily become aware of using it when or even after they do. They simply 'jump to a conclusion', without being conscious of how they got there. Consequently, one can't always know what grounds their type 1 judgements.

For comparison, type 2 thinking, on the other hand, constitutes *deliberate* movements between explicit attitudes in S' 'working memory', that is, in the part of their short-term memory where information is consciously retained and reflected upon (Kahneman, 2003; Stanovich, 2011). For instance, if, instead of making an unconscious, automatic inference, S consciously thought to themselves 'they have let go of the bottle in mid-air', held this proposition before their mind (in their 'working memory'), then added to it 'if they let go of the bottle in mid-air, it will fall to the floor and break', then called to mind the *modus ponens* rule of inference – 'if I believe *p*, and I believe *p* entails *q*, form the belief that *q*' – applied this rule to the first two thoughts, and finally inferred 'the bottle will fall to the floor and break', that would be a very clear case of type 2 thinking[81] (Evans, 2003). Engaging in this kind of conscious reasoning typically requires a lot more time and effort than relying on immediate, automatic, type 1 inferences does. Moreover, because type 2 thinking is deliberate, and takes place in S' conscious working memory, S usually knows what the stages of their reasoning were, and therefore the grounds, or premises, of their conclusion.

Since type 1 phenomena are automatic, S has no agential control over them. This, again, is also true of e-feelings, which lack internal correctability. Thus, e-feelings are regarded as type 1, rather than type 2 phenomena (Proust, 2012). If they're type 1 phenomena, then e-feelings are also 'involuntary', in the sense that I defined in §4.1. To reiterate, here, to say that a state is 'involuntary' means that said state doesn't require anything like S' effort, permission, or agreement to arise: it happens whether S intends or wants it to or not. This is the very definition

[80] Following this rule will normally cause S to form true beliefs, even if it would lead them astray in very atypical circumstances, such as when under the effects of zero gravity.
[81] Again, there's evidence that deductions as simple as *modus ponens* can be performed using automatic, type 1 thinking, so it's important that this example is a case of conscious, deliberate reasoning in S' working memory.

of type 1 phenomena, which are, by nature, automatic. Again, type 1 phenomena are somewhat akin to 'mental reflexes', and happen, when the proper stimulus is applied, without S' effort or permission. For instance, S will form the belief that the glass will fall and break when they see me let go of it in mid-air, whether they want or intend to or not, just as they will raise their leg when their knee is struck by a doctor's reflex hammer, whether they want or intend to or not. Therefore, e-feelings, like perceptual seemings, are involuntary.

Now, some theorists, such as Fischer & Collins (2015) and De Cruz (2014)[82], have argued that philosophical intuitions are type-identical with, specifically, type 1 judgements[83]. Fischer (2018), for example, says:

> "Intuitions are judgments generated by automatic inferences, i.e., by effortless and typically unconscious cognitive processes… When thinkers leap to a conclusion, without being able to explain how or why, prior to sustained reflection, they are relying on automatic inferences, and their conclusions are intuitions" (p.416)[84].

While Fischer and Collins (2015) call identifying philosophical intuitions with type 1 judgements the "psychologically informed aetiological notion of intuition" (p.15), and contrast it against what they deem the traditional, 'uninformed' theories that understand intuitions as the outputs of 'conceptual competence'.

This position differs from my own view, according to which intuitions (at least, intellectual perceptual seemings) are type 1 e-feelings, rather than type 1 judgements. And, at first glance, such a theory isn't obviously absurd, since intuitions do share a number of properties with type 1 judgements. For one, at least some intuitions, as I've argued, are automatic. Moreover, a feature common to most type 1 judgments is, again, that they occur very quickly and easily, and this is likewise true of many intuitions. For example, Cappelen (2012) reviewed a number of seminal thought experiments, such as Burge's arthritis (1979), Thomson's violinist (1971), and Williams' body-swap (1970), and observed that, in each case, the so-called 'intuitions' about those scenarios were relatively effortless to have:

---

[82] Cf. Kahneman and Frederick (2005); Sloman (1996); Evans (2010). A weaker version of the view, which understands some, but not all intuitions as type 1 judgements, also appears in Nagel (2011).

[83] We should charitably assume that, in this case, the word 'judgements' is intended to refer to beliefs, and/or to acts of belief formation, specifically, rather than to type 1 phenomena in general, since it's just too obvious that there are many type 1 phenomena that certainly aren't intuitions, such as FORs relating to metamemory and metaperception, as I'll discuss in the next sub-section.

[84] This wording shows support for a type-, not just token-identity thesis, since it declares both that intuitions are type 1 judgements, and that type 1 judgements are intuitions.

> "[O]ne feature that stands out when these cases are considered: there is some kind of ease, effortlessness, or spontaneity involved. Another way of putting this is that the acts involved don't require a lot of reflection or effort' (Cappelen, 2012, p.33, emphasis in original).

The same intuitions likewise tended to have a very short response time, that is, they arose after only a brief moment of reflection. Indeed, Cappelen suggests equating them, in a deflationary manner, with 'snap judgements'.

Fischer and Collins (2015) have also argued that positing intuitions as type 1 judgements would explain the aforementioned phenomenology of intuitions, that is, the feeling of rightness or correctness that's associated with them. To elaborate, a subject's confidence in their judgement tracks 'answer fluency', or how easy it seemed, introspectively, for them to draw that conclusion (Thompson *et al*., 2011; Simmons and Nelson, 2006; Alter *et al*., 2007; Thompson *et al.*, 2013; Prowse Turner and Thompson, 2009). Again, type 1 thinking is very easy, approaching effortless. Therefore, its outputs are normally accompanied by a strong sense of certainty, or FOR. The aforementioned speed of type 1 judgements likewise contributes here, since subjective confidence also inversely correlates with response time (Kelley and Lindsay, 1993; Robinson *et al*., 1997; Thompson *et al*., 2011; Costermans, *et al.*, 1992). To clarify, FORs aren't essential to type 1 judgements. Indeed, some type 1 judgements will, instead, be accompanied by a FOE, if the subject detects that their automatic response was biased, for instance. However, type 1 judgements, due to their ease and speed, at least tend, very often, to provoke FORs. Therefore, positing intuitions as type 1 judgements might help to explain why they're so strongly associated with such feelings, or so it's been argued.

However – apart from the dubious conflation of intuitions and judgements – there are two problems with thinking that contextualists' intuitions (or intellectual perceptual seemings), specifically, are just type 1 judgements.

First, recall that the intuition that 'scepticism' is true certainly isn't S' immediate impression of that proposition. Rather, they must first take a moment to simulate the sceptic's reasoning themselves, particularly if S is unfamiliar with closure-based sceptical arguments. Indeed, typically, to experience the sceptical intuition, S must reflect upon 'ignorance' and 'closure', hold them in their working memory, and then infer (or, at least, imagining inferring) 'scepticism'. Then, and only then, does S experience the aforementioned 'intuitive pull' of scepticism. Here, a philosophical intuition arises by – indeed, depends upon – performing a conscious, deliberate inference from explicit premises, which is, paradigmatically, thinking of the type 2 variety. After

all, S does, certainly, know what the grounds of their (real or imagined) inference to 'scepticism' was ('ignorance' and 'closure').

Second, type 1 thinking struggles to process negations. As the psychologists Strack and Deutsch (2004) put it:

> "Negating – that is, reversing the truth value of a proposition –… can only be executed in the reflective system" (p.227)[85].

Evidence of this comes from, for instance, the fact that processing negations slows down cognition, suggesting that the, more computationally demanding, and hence much slower, type 2 systems (De Neys, 2006; Stanovich, 1999) are necessary to comprehend negated concepts (Wason, 1959). Moreover, because type 2 processes generally take longer to complete, the amount of time allocated to a cognitive task influences the extent of type 2 engagement with it, with less time correlating with more reliance on fast, type 1 strategies instead (Evans and Curtis-Holmes, 2005; Finucane *et al.*, 2000; De Neys, 2006; Roberts and Newton, 2001; Tsujii & Watanabe, 2010). For example, if asked to respond quickly, experimental participants are more likely to answer on the basis of type 1 believability biases (also known as 'conclusion-believability' heuristics) (Evans and Curtis-Holmes, 2005; Tsujii & Watanabe, 2010). Several studies attest to the fact that negations are instead processed as affirmations when thinking time is limited, thus suggesting that parsing negations *as* negations is beyond the scope of type 1, and requires the use of slower, type 2 thinking instead (Gilbert *et al.*, 1990; Gilbert *et al.*, 1993).

Negation appears in every premise of AI. 'Ignorance' negates the BIV scenario, while both 'closure' and 'scepticism' negate the concept of knowledge, by stating that S *doesn't* know, for example that they have hands. Consequently, we would expect these propositions, under normal circumstances, to be parsed by type 2, not type 1, thinking, and thus for any thoughts that S has about them to be (at least partially) outputs of the former.

So, for two reasons, contextualists' intuitions about AI arise from, and indeed require, some degree of type 2 thinking. Consequently, we can't simply equate (these kind of) philosophical intuitions with purely type 1 phenomena, such as judgements arising solely from type 1 thought.

Positing intuitions as e-feelings, specifically, circumvents these issues. To elaborate, while type 1 thinking is more likely to trigger FORs, for example, it isn't performing type 1 thinking *per se* that causes this feeling in S, but rather them having, for instance, a fast and fluent thought. Therefore,

[85] To clarify, this doesn't mean that type 1 thinking can't *conclude* that something isn't the case, or is false; it's specifically parsing negations as negations that it struggles with.

while e-feelings themselves are type 1 phenomena, they can be triggered by either type 1 or type 2 thoughts, if those thoughts are sufficiently fluent, for example. For instance, if the explicit, type 2 inference from 'ignorance' and 'closure' to 'scepticism' is sufficiently easy for S to perform, or simulate, in their working memory, then AI will feel sound to them; here, it's properties of type 2 reflection, specifically, that triggers e-feelings. Thus, unlike the position that philosophical intuitions are just type 1 judgements, my view, that said intuitions are e-feelings, allows for some intuitions to arise from type 2 thinking. Put simply, e-feelings can arise from type 2 thinking, just as contextualists' intuitions, at least, do.

So, my view refines the burgeoning type 1 thinking account of intuition: rather than construing intuitions as type-identical with type 1 judgements, we should picture them as token-identical with certain e-feelings, which are likewise type 1 phenomena, but which can be triggered by either type 1 or type 2 thinking.

Non-Inferential. E-feelings, like perceptual seemings, are non-inferential.

Another feature of type 1 phenomena, and thus of e-feelings, is that they don't arise from conscious inferences. Rather, they're the (conscious) outputs of entirely unconscious cognitive processes. For one, as I explained earlier, S doesn't usually know what the grounds of their type 1 thoughts is. For example, turning to e-feelings, specifically, when S has a particularly fluent thought, and experiences a FOR toward its content, they don't normally know, or sense, introspectively, that their FOR was grounded in (caused, or explained by) said fluency. Indeed, the fact that FORs are often triggered by fluency took considerable *a posteriori* psychological research to uncover. In fact, introspectively, type 1 thoughts don't seem to be based on anything else, 'lower down', at all (Fischer and Collins, 2015). To put it simply, our introspective access to the grounds of our reasoning 'bottoms out' at our type 1 thoughts. However, if such type 1 thoughts were the result of conscious inferences between explicit propositions, then S would know, or at least have introspective access to, the grounds, or premises, of their conclusions. Therefore, type 1 thoughts, and thus e-feelings, are, like perceptual seemings, non-inferential, in the sense that I defined in section §4.1.

Gradable. E-feelings, like perceptual seemings, come in degrees of vividity, or exactness.

E-feelings can have more or less vivid, or exact, content. Consider, for instance, the 'feeling of forgetting' (or FOF) (Arango-Muñoz, 2014). This is the feeling that S gets when, for example, they're about to leave the house, and sense that they've forgotten something important, though they cannot recall specifically what. The content of a FOF – what S feels they've forgotten – can range from an extremely vague, nonspecific 'something', to the, slightly narrower, 'to bring

something' or 'to do something', to the, more exact, 'to bring something to do with that reading group… begins with a T…', and so on. Thus, e-feelings, like perceptual seemings, come in degrees of vividity, or exactness.

Tempting. E-feelings, like perceptual seemings, tempt belief.

Finally, as I've already explained, in §5.1.2, one function that e-feelings have is to motivate belief or disbelief in their content. For instance, to reiterate, if S experiences a FOR or a FOE toward *p*, this can, subsequently, make them feel relatively more motivated to accept or reject *p*, respectively. Thus, e-feelings, like perceptual seemings, can give rise to doxastic temptations.

So, to summarise, e-feelings bear the same ten resemblances to perceptual seemings as contextualists' intuitions did. Therefore, we should think of e-feelings as, likewise, a variety of perceptual seeming.

In the next sub-section, I'll argue that certain FORs and FOEs have an intellectual aetiology, and are, therefore, intellectual perceptual seemings, specifically. I'll thus conclude that certain philosophical intuitions, namely intellectual perceptual seemings, are type-identical with this variety of e-feelings.

5.2.2 Intellectual FORs and FOEs

Some e-feelings clearly aren't 'intuitions', in the sense that contextualists use the term. When S is trying to remember someone's name, for example, and experiences the TOT feeling, or is asked a question, and experiences the FOK, they're not having an 'intuition', or 'intuiting', in the relevant sense, nor are they doing so when they're merely surprised or curious. This is because a contextualist's intuition is, again, a feeling that a proposition, like 'scepticism', is true (or false), or likely (or unlikely), or that an argument, like AI, is valid (or invalid), or sound (or unsound). By contrast, the TOT and FOK are feelings about the probability that a present or future mental action, respectively, will succeed, while surprise and curiosity are more to do with whether an item of information conflicts with S' prior expectations, or invites further investigation, than with the accuracy or soundness of said information. To put it simply, some e-feelings have different content (actions, rather than propositions or arguments), or ascribe different epistemic evaluations (probability of success, rather than accuracy or soundness), from contextualists' intuitions. Another way to put this is to say that said e-feelings are triggered by the wrong

conditions (thinking[86] about actions, rather than propositions or arguments, or thinking about whether *p* is expected or interesting, rather than whether it's true) – or, in other words, have the wrong aetiology – to be the relevant kind of philosophical intuitions.

Some FORs and the FOEs, however, fit the aetiological profile of contextualists' intuitions much more neatly. For example, when S just thinks, in the strong or weak sense, a proposition, *p*, and experiences the feeling that *p* is true, S is experiencing the very feeling described by contextualists using intuition-talk, namely, that a proposition is true when they consider its accuracy. We can call said e-feelings the 'intellectual' FORs and FOEs, being FORs and FOEs that arise, specifically, when S merely thinks about some proposition, and it feels to them that it's true (or false), or probable (or improbable), or performs or simulates some inference or argument, and it feels to them that it's or valid (or invalid), or sound (or unsound). In other words, these are FORs and FOEs triggered by 'armchair reflection', broadly construed, upon claims and arguments.

So, e-feelings are perceptual seemings. Moreover, some have an intellectual aetiology, and are therefore intellectual perceptual seemings, specifically.

In conclusion, both contextualists' intuitions and certain e-feelings are intellectual perceptual seemings. Therefore, some philosophical intuitions are a certain variety of e-feelings.

It is worth clarifying that this view should not come across as in any way uncharitable to those philosophers, such as the contextualists, who appeal to such intuitions as evidence, neither in the sense that it burdens them with systematic error about their own intuitions, nor in the sense that it attributes to them an obviously flawed philosophical method.

It is not an uncharitable theory in the former sense because, while the contradictions in the ways that contextualists describe their intuitions mean that I cannot produce a coherent theory of intuition that validates *everything* they say about them, my theory does respect their self-reporting as much as humanly possible. In fact, one might say that the whole aim of this half of the thesis has been to arrive at a theory of intuitions that, unlike extant views, doesn't blatantly contradict the way that contextualists talk about their intuitions. Looking back, you'll see that my overall approach has quite simply been to assume that the contextualists are generally reliable and honest when it comes to reporting their own mental states, and derive the properties of their intuitions from how they most commonly choose to describe them. Conversely, it is *other* theories of intuition that would burden them with systematic error. For example, contextualists describe

[86] 'Thinking' here obviously need not mean consciously reflecting upon, since e-feelings can be triggered by entirely unconscious thought.

their intuitions as ‘feelings’ much more often than they ever describe them as speech acts, and describe them as things that they disagree with – or otherwise fail to believe – a lot more often than they describe them using words like ‘belief’, so doxastic or ordinary language accounts *would* burden contextualists with systematic error about their own minds, and thereby be uncharitable. My own theory, by contrast, only attributes to them minor and infrequent stylistic errors in the manner that they choose to describe their intuitions in prose. I’ve even included explanations (or ‘error theories’) throughout my work for why philosophers might be tempted to employ such, ultimately misleading, descriptions on occasion. Additionally, it isn’t uncharitable to say that philosophers don’t know when their intuitions are epistemic feelings. E-feelings are a recent topic of academic investigation, and have thus far been almost exclusively within the purview of psychologists, with only very minimal, if steadily growing mention of them appearing in philosophy. Therefore, it’s entirely unsurprising that contextualists, for example, don’t know of these phenomena, or know enough about them to introspectively equate them with the mental states of theirs that they’ve been calling ‘intuitions’.

My theory isn’t uncharitable in the latter sense either, for three reasons. First, because whether e-feelings are a reliable source of evidence is an open question. It certainly isn’t the case that such feelings are obviously or *prima facie* bad at leading us to the truth. Indeed, there’s at least some initial reason to think that they’re actually quite trustworthy, especially in individuals with high cognitive capacity, such as those typically found within academic circles (Thompson, 2014). Now, one might think that, if e-feelings do, in some future time, turn out to be poor evidence, then my thesis is dubious, because philosophers – who are intelligent, conscientious individuals – wouldn't base their theories on something so unreliable. However, while it is uncharitable to attribute to philosophers a methodology that is obviously or likely flawed, it is not uncharitable to attribute to them a methodology that seems superficially, at this time, to be valid, but which might, one day, following future research, turn out to be defective. We academics are constantly learning how to refine and improve our research and argumentative methods, and it would be strange to assume that modern thinkers have arrived at the end of this process, or call it uncharitable to assume otherwise.

Second, even if it does turn out that relying evidentially on e-feelings is a flawed approach to philosophy, there’s an obvious explanation, or 'error theory', for how philosophers could have made the mistake of trusting them to begin with. We’ve seen that e-feelings give rise to doxastic temptations, inclining their subjects to accept their content. In other words, the phenomenology of e-feelings in a sense presents them (or their content) as trustworthy, or at least makes it difficult for people *not* to trust them. That philosophers have ‘given in’ to trusting a mental state

the function of which is, at least partially, to make subjects accept its content (for the purposes of metacognitive control and correction) is hardly an egregious error.

Third, and finally, it should be noted that there are limits to interpretive charity. It's clear enough from the vague, amorphous way that contextualists describe these mental states that they haven't given nearly enough thought to what their intuitions actually are. I would not be at all surprised if the recent surge in metaphilosophical research unearthed some pretty inconvenient truths about the methods being practiced by such individuals: if, when we do look more closely at this phenomenon, it turns out to be an unreliable guide to philosophical truth. That is, I don't take it to be that unlikely at all that the methods being practiced by some, less methodologically reflective philosophers are defective.

So, to reiterate, I've argued that some philosophers, such as contextualists, use intuition-talk to refer to certain kinds of epistemic feelings. Thus, two, hitherto largely segregated academic domains – the philosophy of intuition and the study of e-feelings – are poised to mutually inform one another, to the advantage of both. In Chapter 6, I'll briefly consider what insights this reciprocal relationship might yield, and some further implications for identifying philosophical intuitions with e-feelings. First, however, I'll address an anticipated objection to my account.

5.3 Epistemic Feelings and Metarepresentation

In this section, I anticipate a significant objection to my account of (a certain kind of) philosophical intuitions – namely, that it commits me to saying that all said intuitions are metarepresentational – and offer a comprehensive response.

As I explained in §5.1, e-feelings are often assumed to be metacognitive. However, there are many lengthy and heated debates about metacognition in the literature. One such debate takes the form of a disagreement over whether metacognition requires metarepresentation. Metarepresentation is when one of S' mental states explicitly represents mental states, and/or facts about them, *as* mental states, and/or facts about them. For example, when S believes that 'S is having a perceptual experience of an apple on the table', said belief metarepresents; specifically, it represents a perceptual experience *as* being a perceptual experience, and a fact about that experience, that its content is an apple, *as* being a fact about said mental state.

The view that all metacognition involves metarepresentation is the view that all (genuine) metacognition is what's sometimes called 'high-level' metacognition (Carruthers, 2009; Flavell, 1998). High-level metacognition is self-regulation based on metarepresentational attitudes. For instance, another example of a metarepresentational state would be the belief that 'S' memory is

unreliable when it comes to names', and an instance of high-level metacognition is S exercising lower confidence toward their apparent recollection of names on the grounds of that metarepresentational attitude. This is also known as 'theory-based' self-regulation (Koriat and Levy-Sadot, 1999), since metarepresentation requires possession of folk-psychological concepts, such as 'memory', and thus at least a primitive 'theory of mind' (Proust, 2012, p.241).

Some, however, also believe in what's sometimes called 'low-level' metacognition (Arango-Muñoz, 2011, 2014, Endnote 14; Koriat, 2000; Proust, 2007, 2009a, 2009b). Low-level metacognition is, by contrast, self-regulation that isn't based on metarepresentational attitudes (Schwarz, 2002), and thus doesn't require a theory of mind. On the contrary, low-level metacognitive processes detect and respond to features of S' first-order cognition, and provide feedback on it, without those monitoring and control mechanisms, or their outputs, explicitly representing S' first-order states and processes, and/or facts about them, *as* mental states and processes, and/or facts about them. For instance, while theory-based (or 'theoretical') metacognition might, in response to detected features of S' belief that *p*, output a metarepresentational attitude that represents S' belief that *p as* a false belief with content *p*, low-level metacognition (which Proust sometimes calls 'procedural' metacognition) might, in response to detected features of S' belief that *p*, output a non-metarepresentational attitude that only has the, first-order content that *p* is false, not that *p* is believed, and/or that said belief is inaccurate (Proust, 2012, p.251).

There are corresponding disagreements over whether e-feelings should be considered metacognitive after all. If they're not metarepresentational, then those who deny the existence of low-level self-regulation, and/or that said regulation constitutes genuine metacognition, will deny that they're truly metacognitive (Carruthers, 2017). Meanwhile, those who wish to include e-feelings among metacognitive phenomena either argue that low-level self-regulation is possible, and constitutes genuine metacognition (Proust, 2012) or that e-feelings are metarepresentational (Nagel, 2014).

Nothing I've said in this thesis hinges on whether e-feelings are genuinely metacognitive. However, those who argue that e-feelings are metarepresentational might take issue with my view that (some kinds of) philosophical intuitions are e-feelings. Specifically, they might argue that said intuitions aren't always about, don't always 'metarepresent', mental states or processes. When S intuits that 'ignorance' is true, for instance, S typically intuits only that the proposition is true, not that their belief in 'ignorance', say, is correct. Indeed, contextualists, at least, never once describe the relevant kind of intuitions as being about, or 'metarepresenting' their own mental states; again, such intuitions are about first-order propositions and arguments, not mental states

or phenomena. Moreover, in §4.6, I argued against the view that said intuitions are doxastic temptations, claiming that they represent propositions, rather than the mental acts of judging them, essentially saying that contextualists' intuitions have first-order, not metarepresentational content.

In this section, I respond to this anticipated objection by arguing that e-feelings aren't, invariably, metarepresentational. First, in §5.3.1, I'll extract two existing arguments from the wider, aforementioned metacognition debate in favour of the view that e-feelings are metarepresentational, and offer some responses. Then, in §5.3.2, I'll extract and expand upon two existing arguments from the wider metacognition debate against the view that e-feelings are metarepresentational.

5.3.1 Arguments for Metarepresentation, with Rebuttals

First, it's been suggested that the evaluative component of e-feelings makes them metarepresentational. Specifically, that because e-feelings convey (epistemic) information about past, present, or future cognitive processes to their subject, they represent those processes, and facts about them. Some, for example, have claimed that surprise is explicitly metarepresentational: that it signals to S that their experiences are in conflict with their prior expectations or beliefs, and therefore represents S' experiences to S (together with the epistemic assessment that they contradict some prior expectation or belief) (Davidson, 1982).

The present understanding, however, is that feelings of surprise are merely triggered by such conflicts, but don't represent them, or their constituent mental phenomena (Carruthers, 2017, p.66). For one, consider that S might feel surprised without knowing, immediately, what exactly is surprising to them about an experience, and/or why, and have to reflect upon it, and/or upon their own expectations and beliefs, to find out. In other words, S' cognitive systems might detect a conflict in their mental life, and signal to S, in the form of an e-feeling, that such a conflict is present, without that feeling, in itself, conveying to S any particular experience, or any specific belief or expectation that it contradicts. Thus, the feeling of surprise need not embed a representation of S' experiences, beliefs, or expectations. Indeed, we might speculate that surprise exists, at least in part, to motivate S to investigate their experiences and expectations to discover the cause of the conflict. That is, surprise could just be an initial sensation triggered by the detection of a nonspecific contradiction (by the activation of processing pathways that interfere with each other, for instance), that draws S' attention to their present experiences, and compels further investigation.

Moreover, consider that human infants are typically thought, by developmental psychologists, to experience surprise – evidenced by, for instance, their longer eye contact with unexpected objects (see, for example, Yeung *et al.*, 2016) – but not to possess the mental concepts necessary to represent their mental states *as* mental states (Bermùdez, 2009; Carruthers, 2008); that is, they're capable of surprise, but not yet capable of metarepresentation. Essentially, it's thought that very young children (Balcomb and Gerken, 2008), as with most non-human animals (Hampton, 2001; Smith, 2009), are still 'sensitive' to their cognitive processes, or facts about them, even if they can't conceptualise them *as* cognitive processes, or facts about them. In other words, their cognitive monitoring and control mechanisms can still produce surprise, in response to a contradiction in their mental lives, and that feeling can still directly affect their subsequent thoughts and behaviour, even if they cannot conceptualise the mental constituents of the contradiction *as* mental phenomena. To put it simply, surprise doesn't require metarepresentation.

Second, it could be suggested that the functions of e-feelings make them metarepresentational. Again, surprise, for example, motivates S to investigate their experiences more closely, or for longer, and/or their own expectations and beliefs, to help discern what, more exactly, is surprising about a particular state of affairs, or which of S' prior beliefs or expectations it's in tension with, or to confirm whether the apparent contradiction is genuine, or to help them make sense of the anomalous data, and so on. Here, one might say that surprise signals to S that they should pay closer attention to an 'experience', acquire more 'beliefs' about it, introspect upon their own 'expectations', and so on. Therefore, one might argue that, even if the evaluative component of e-feelings doesn't represent S' mental states, or facts about them, such feelings do still represent mental states, specifically as a *goal* for S to work towards; that is, e-feelings represent mental states as a desirable or worthwhile end.

This second line of argument is vulnerable to the same criticism that infants can experience surprise, and be motivated to, say, look longer at an unexpected object, even if they lack mental concepts, such as 'look' or 'visually experience'.

Moreover, as Carruthers (2017) likewise observes, it isn't necessary for e-feelings to represent any goal, let alone a mental one, to motivate behaviour. Indeed, it's been observed that folk psychology often overintellectualises human behaviour by describing its motivation in terms of goals; that is, we prefer to talk as though all actions are 'goal-directed' (Milner and Goodale, 1995), overlooking the fact that many of them are 'directly' motivated by our feelings (Panksepp, 1998). Suppose, for example, that S sees a bear, experiences fear, quickly runs away from it, and is later asked 'why did you run away?'. S might be tempted to give an answer like 'because I

wanted to stay alive', thereby citing some goal – staying alive – as part of what motivated their retreat. However, on reflection, it overcomplicates the phenomenon of fear to say that, when we're afraid, our running away is triggered by some particular objective that we have in mind. A more accurate explanation of the retreat would, at many times, be much simpler: just 'because I was scared'. That is, fear can directly motivate retreat, without bringing to mind a specific end. The feeling of fear compels withdrawal from its cause, or object, even without representing any further goal.

Likewise, surprise can, for instance, motivate further investigation of a phenomenon in a similarly direct, rather than goal-oriented, manner. S doesn't approach a surprising object, for example, with the aim in mind of acquiring 'beliefs' about it, or 'experiencing' it for longer, rather, they're simply drawn to it by the phenomenology of the e-feeling. As Carruthers (2017) puts it:

> "[C]ommon-sense psychology seems to require one to cite a metarepresentational goal when explaining such behavior. ("I *wanted to know* what it was, so I walked over to look.") But in reality, curiosity is a first-order affective state, which includes a first-order motivating question-like state, together with a suite of primed information-gathering behaviors" (p.68, emphasis in original).

In short, e-feelings can compel mental actions, such as acquiring beliefs or having experiences, without representing them.

Indeed, a state's 'mental affordances' doesn't entail anything about its representational content. For example, if S has a visual experience as-of there being an apple on the table, this will normally afford (meaning cause or motivate) S to form a corresponding belief that there's an apple on the table. However, said visual experience certainly isn't about S' beliefs but is, rather, about the world. Therefore, even if a mental state affords mental states, this doesn't entail that it's about, or represents, such states. Likewise, although surprise motivates S to acquire beliefs, this doesn't entail that S' surprise is about beliefs. Indeed, nearly all mental states afford other mental states or actions (often in addition to, rather than instead of, non-mental actions): they influence other, 'downstream' cognitive processes (Carruthers, 2017, p.67). However, it would be highly controversial to claim that nearly all mental states are metarepresentational: that they're all 'about' the other mental states or actions that they afford.

So, superficial reasons to think that e-feelings might be metarepresentational are, on reflection, unconvincing. In addition, there are at least two positive arguments against e-feelings being metarepresentational.

5.3.2 Arguments Against Metarepresentation

First, Proust (2014) argues that e-feelings "neither refer to all mental states nor to a particular one" (p.144), because, while the feeling of fluency, for example, can arise in reaction to a wide range of mental activities – remembering, perceiving, judging, and so on – it fails to convey which of these activities, specifically, triggered it. In other words, the feeling conveys no information about *which* mental process is fluent and, therefore, doesn't represent said process. Her evidence for this comes in the form of some well-documented cognitive confusions. For example, when people are asked to make 'judgements of learning' – to assess whether they've 'learned', and will therefore be able to recall, at a future time, a specific piece of information – they'll find information that they're currently very familiar with (say, from having just read it) fluent (fast and easy to parse), and then misinterpret the fluency of processing the information as a predictor of the fluency (speed and ease) with which a future recall of that data will occur, causing them to overestimate their retention. In other words, the fluency that S experiences doesn't testify, specifically, that their current processing is fluent, rather than their future recall. That is, fluency doesn't represent a specific cognitive activity. Indeed, I'll add to this that the evidence here supports the view that the information itself, *p*, is what feels fluent to S, rather than any processing or recollection of it. This would explain why said fluency makes S confident in any cognitive activity involving *p*, rather than in a specific mental task, but would also make fluency an entirely first-order, non-metarepresentational experience.

The only issue I have with Proust's argument is that fluency itself isn't, uncontroversially, an e-feeling. To elaborate, as we've seen, the feeling of fluency can give rise to either a FOR or a FOE, depending on whether S associates it with accuracy or error. This suggests that the feeling of fluency doesn't, in itself, convey any particular epistemic assessment of its object. After all, an item of information can feel just as fluent regardless of whether S assesses it positively or negatively. In other words, fluency grounds, but isn't identical with, e-feelings. Nonetheless, the data that Proust cites here also implies that the FOR, for example, is likewise non-metarepresentational. To elaborate, it shows that, when occurrently processing *p* feels fluent to S, this fluency grounds a FOR, or a feeling of confidence, not just toward their occurrent processing, but also toward their future recall process. So, again, their confidence extends to *p* itself, and thus, liberally, to any task involving *p*, in general. Therefore, just as S' feeling of fluency is directed to *p* itself, rather than to any specific cognitive task involving *p*, their FOR is experienced toward *p* itself, and is likewise non-metarepresentational.

Another objection to Proust comes from Nagel (2014), who argues that feelings of fluency – or, again we should say, FORs – might still always be about a specific cognitive process, even if S can

be wrong about which one, and mistakenly believe themselves, for instance, to be feeling confident about a future recall attempt, rather than about their occurrent processing. While this is technically compatible with the evidence, it's a somewhat *ad hoc* way of defending the metarepresentational interpretation of e-feelings. Unless independent reason is provided to think that subjects really are systematically mistaken about the metarepresentational content of their e-feelings, I think it safe to assume that such feelings simply fail to convey any metarepresentational information at all.

Second, Proust (2012) argues that e-feelings don't metarepresent because they present the subject with solutions to the task that they're trying to perform, or with the available options that S has for achieving it, and thus guide them toward the desired outcome, rather than representing their first-order cognitive processes that are attempting it, or facts about them. Here, Proust seems to agree, broadly, with Carruthers and Ritchie (2012) that e-feelings are:

> "directed at the world (in particular, at the primary options for action that are open to one), rather than at one's own mental states" (p.82)[87].

Proust's argument for this is that, because e-feelings guide S' cognition, they signal what she calls 'knowledge affordances', that is, they tell S how to go about achieving their desired outcome, which need not require attributions of mental states to S:

> "the reason why [metarepresentation] is not needed for feelings to guide decisions is that this guidance is procedural… thoughts expresses the specific knowledge affordance associated with a first-order cognitive performance, and thereby guides epistemic decisions" (Proust, 2014, p.144).

The FOR, for instance, presents, say, a proposition, *p*, as a viable answer to a question, rather than representing the cognitive faculties that are endeavouring to answer said question. That is, the FOR serves to highlight and motivate specific answers, or solutions, to mental tasks, rather than representing the mental processes that are undertaking it, or facts about them. Certainly, representing such mental phenomena isn't necessary for motivating S toward specific (mental or physical) actions, and thus guiding their first-order behaviour.

[87] For transparency, I'll note that, for Carruthers, this lack of metarepresentation is sufficient to render e-feelings non-metacognitive, whereas Proust (2012) believes in a form of non-metarepresentational procedural metacognition, grounded in the "functional coupling" (p.250) between e-feelings and the processes they govern. Again, my concern in this section is only with metarepresentation, not metacognition.

Nagel (2014) has since objected to Proust that there's no tension between metarepresentation and procedural guidance, since a metarepresentational state that attributes mental phenomena, or facts about them, can likewise guide behaviour. For instance, there's no reason why S' behaviour, such as their reliance on their apparent recollections, couldn't be guided by, for example, the metarepresentational thought 'my recollection of names is unreliable'.

However, even if we acknowledge that metarepresentational states can likewise guide behaviour, this doesn't entail that S' behaviour isn't also, at times, guided by non-metarepresentational e-feelings in the manner that Proust describes. Indeed, picturing all procedural guidance as metarepresentational would certainly overintellectualise. When S is looking for their car keys, for example, and relies on their memory to help guide the search, their memorial faculties deliver the last few places that S remembers seeing or having them. That is, S' memory directly guides their actions by offering up first-order answers to the question 'where are my keys?' – answers like 'in your coat pocket' or 'in the top drawer of your desk' – motivating them to search those areas. What S' memory doesn't do is indirectly guide the search by representing their memories to them *as* memories – conveying metarepresentational information like 'S has a memory with content *p*' – which they subsequently take under advisement when deciding where to look. In other words, S' behaviour is guided by representations of worldly places and locations, rather than by representations of their own mental states, or facts about them. There's simply no reason to think that the procedural guidance afforded by e-feelings, specifically, should be much different: rather than presenting facts about S' own cognitive activity to them, on which grounds they can make decisions about how to reason, it likely just offers up possible first-order answers – say, '*p*' – when S is trying to, for instance, arrive at true beliefs about a matter. At the very least, the onus of proof is squarely on Nagel to show that the function of e-feelings is different in kind from other forms of procedural guidance.

In conclusion, e-feelings aren't, on reflection, invariably metarepresentational, and thus my positive account of intuitions can easily be reconciled with their non-metarepresentational content.

5.4 Chapter 5 Conclusion

In conclusion, certain philosophical intuitions are intellectual perceptual seemings, and intellectual perceptual seemings are a certain class of epistemic feelings, specifically, a variety of FORs and FOEs. Thus, some philosophical intuitions, at least, are epistemic feelings. Since said intuitions are, as I've shown, the primary support for EC, such epistemic feelings have, through the burgeoning influence of EC, played a significant part in shaping contemporary epistemology.

In the next, and final, chapter, I summarise my findings in this thesis, and briefly consider some of their possible implications.

Chapter 6 Final Conclusion

It's widely assumed that intuitions are central to the methods of contemporary analytic philosophy. In particular, it's thought that philosophers appeal to intuitions as evidence, or as a source of evidence, for their claims. Recently, this view of philosophical practice has become known, in some circles, as 'centrality'.

In this thesis, I addressed two questions about centrality, namely:

A) Is centrality actually true? and
B) What are intuitions?

As a partial answer to question A, I argued that intuitions do, indeed, play a central role in, at least, some core philosophical arguments, and helped to clarify that role. Specifically, I made the case for a centralist interpretation of the primary argument for epistemic contextualism, identifying, through a conscientious analysis of the most seminal literature, not one but several, specific ways in which contextualists appeal to intuitions in an evidential capacity. Moreover, since contextualism is chiefly motivated by said argument, and is having a burgeoning influence on modern epistemology, and considering that epistemology is of ubiquitous philosophical significance, with ties to arguably all other core philosophical topics, such as ethics and metaphysics, I thereby demonstrated that treating said intuitions as evidence is profoundly shaping the discipline at large.

Then, as a partial answer to question B, I developed a novel account of what some philosophical intuitions – namely, those of the kind cited in the aforementioned argument for epistemic contextualism – are. In Chapter 3, I argued against the extant 'minimalist' views that such intuitions are outright beliefs, partial beliefs, dispositions to believe, or facts about ordinary language. I also argued, in Chapter 4, that existing accounts of intuitions as 'intellectual seemings' are too conservative to capture the intuitions in question. In place of these alternatives, I defended a more liberal version of the intellectual seemings thesis. Then, in Chapter 5, I argued that said seemings, and hence certain intuitions, are a sub-category of what some psychologists call 'epistemic feelings'. Epistemic feelings are experiences triggered by metacognitive monitoring and control subsystems, in response to features of a first-order cognitive process and/or its outputs, such as its fluency, and exist to regulate said first-order cognition, such as by encouraging or dissuading belief, or drawing the subject's attention to contradictions between assimilated information and their own prior expectations. More precisely, I argued that the relevant intuitions are 'feelings of rightness' (or FORs) and 'feelings of error' (or FOEs) with their source in thought,

or reflection. This interdisciplinary thesis revolutionises our understanding of (some kinds of) philosophical intuitions, and bridges two, hitherto largely segregated academic sub-disciplines: the philosophy of intuition and the study of metacognition.

I therefore conclude, overall, that metacognitive experiences have profoundly shaped, and by all accounts are continuing to shape, modern philosophy.

In this, final chapter, I briefly consider some of the implications of my discoveries (§6.1). Then, I outline what future, productive research this invites (§6.2).

6.1 Summary of Findings and their Likely Implications

In Chapter 1, I demonstrated that philosophical practice is hugely varied, both in terms of what philosophers mean by words like 'intuition', and in terms of how they employ said intuitions in their arguments. Consequently, metaphilosophers should avoid overgeneralising, since the nebulous nature of philosophy renders any general claims about it – how it proceeds, and whether those methods are reliable – vulnerable to counterexamples.

Indeed, as promised, throughout the thesis we've witnessed examples of how much existing metaphilosophy does overgeneralise, making excessively broad claims about the discipline as a whole, to its downfall. For example, the arguments against centrality that I discussed in Chapter 2 all concluded that intuitions, (overly) broadly construed, aren't central to analytic philosophy, (overly) broadly construed, and thus ran afoul of cases where a *certain* sense of intuition is used as evidence in a *certain* way, or ways, and where this practice really does shape the core of philosophy. Likewise, the theories of intuition that I discussed in chapters 3 and 4 were claims about the general meaning of 'intuition' in philosophy that failed to specify any particular sense of intuition, or area of philosophical investigation – any more narrowly than 'conceptual analysis' in general – to which they apply, and thus, again, ran into counterexamples: cases where intuition-talk clearly *doesn't* mean what they claim it to.

On the basis of these concerns about overgeneralisation, I made the case for a much more focussed, modest approach to metaphilosophical investigations. This, I argue, is the way in which all future metaphilosophy should proceed.

In keeping with this, I should, and will, be equally modest about the implications of my own findings. Accordingly, I will, for the most part, limit the scope of said implications to only the very narrow domain of literature that I've reviewed herein, and be extremely cautious about drawing more general conclusions about either first- or higher-order philosophy as a whole.

In Chapter 2, I showed that the argument for epistemic contextualism appeals to intuitions in an evidential capacity and, indeed, in multiple senses: intuitions are treated as data to be explained by a theory, as a reliable guide to truth, as a means of unearthing our 'Socratic' knowledge, as a way to pinpoint their content, and as justification to believe their content. To reiterate, since contextualism is chiefly motivated by said argument, and is having a burgeoning influence on modern epistemology, and considering that epistemology is of ubiquitous philosophical significance, I thereby demonstrated that treating said intuitions as evidence is profoundly shaping the discipline at large. This finding, at least, does have wider implications for the centrality/anti-centrality debate. To elaborate, while centrality might still, potentially, be false, insofar as the majority of philosophers don't appeal to intuitions as evidence, the anti-centralists are, either way, mistaken in thinking that intuitions haven't been, and aren't continuing to be, central to the evolution of modern philosophy. In other words, as I explained in Chapter 1, the 'spirit' of centrality, at least, is true, and so the study of intuitions could indeed have profound metaphilosophical implications (*contra*, for instance, the aforementioned views of Cappelen). This, more than anything, justifies the branch of research that dominated the second half of this thesis: the investigation into *what* these 'intuitions' actually are.

Moreover, my findings help to clarify what the role of intuitions in philosophy looks like in much more detail than the general, broad characterisation of centrality that we started out with, namely:

> *Centrality*: Contemporary analytic philosophers rely on intuitions as evidence (or as a source of evidence) for philosophical theories.

That is, I've described several, more specific ways in which centrality is practiced. Again, this isn't intended as a general claim about what the reliance on intuitions in philosophy, as a whole, looks like. Rather, this is just beginning to work out some of the specific practices that are out there.

Moreover, while treating intuitions as data to be explained, as a reliable guide to truth, and as justification for belief, are perhaps more obvious and widely recognised possibilities, the 'Socratic' and 'introspective' kinds, at least, are newly identified forms of centrality, and novel additions to the taxonomy of philosophical method. While evaluative metaphilosophical research has been done on how to best explain intuitions, whether they're reliable, or can justify, little to no attention has been heeded to these other practices, which invite their own assessment.

In Chapter 3, in my destructive project, I showed that contextualists' intuitions aren't outright beliefs, partial beliefs, dispositions to believe, or facts about ordinary language. This has significant implications for future metaphilosophy. Narrowly, it means that researchers looking

into contextualism, specifically, must be extremely cautious about how they conceptualise the relevant intuitions. We cannot straightforwardly determine, for example, whether laypeople have these alleged intuitions merely by asking experimental participants to report their beliefs, or degrees of agreement, or by recording normal patterns of speech behaviour, nor can we assess the reliability (or the evidential merits more generally) of said intuitions just by looking at how judgements and credences react with truth-irrelevant parameters. To reiterate, while we might expect people's intuitions to normally track their beliefs, credences, and speech acts, it's exactly in philosophically interesting cases, which give rise to paradoxes, such as those which are the focus of contextualism, where we should most often expect these phenomena to come apart. In other words, this thesis provides some guidance for how, or how not, to perform experimental work on the intuitive basis for contextualism: it tells us what kind of phenomenon to *not* look for and epistemically assess in our metaphilosophical critique of the case for EC.

More broadly, meanwhile, my findings suggest that researchers need to be very careful about what assumptions they make regarding what intuitions actually are and how they, or their accuracy, can be measured, whether they're looking at contextualism, specifically, or not. One cannot, for instance, simply assume that intuitions are, for instance, beliefs (of some kind), since, demonstrably, there are at least some, important areas of philosophy wherein they aren't. For instance, while it is certainly possible that other philosophers, arguing in other areas for different theories, are, unlike these contextualists, appealing to beliefs, the possibility that they are not doing so should at least be considered. At the very least, before undertaking any experimental work on what intuitions people have, and whether this accords with philosophers' empirical-psychological claims, for example, a conscientious investigation of the target literature must first be conducted, to clarify what the relevant phenomenon, 'intuitions', actually is.

In Chapter 4, at the beginning of my constructive project, I defended an interpretation of contextualists' intuition-talk as referring to intellectual perceptual seemings. I argued against extant intellectual seemings theories as an account of these intuitions, specifically, on the grounds that they're overly conservative, and defended a more liberal position instead. I also showed how, more generally, intellectual seemings accounts of philosophical intuitions can, and indeed should be reconciled with a presentational, or quasi-perceptual model. Finally, I argued that an intellectual perceptual seeming isn't a doxastic temptation, or feeling tempted to form a belief, but rather just a proposition feeling true. Overall, this chapter provided a more developed version of the intellectual seemings thesis than has previously been put forward in the literature. Indeed, in this thesis I've identified no fewer than ten properties of the relevant seemings, explicated their presentational structure, and explored their phenomenology. This furnishes future researchers with a much more thorough picture of what some intuitions, at least, actually look like.

Finally, in Chapter 5, I argued that contextualists' intuitions, or 'intellectual perceptual seemings', are actually a sub-category of epistemic feelings, specifically, a certain variety of 'feelings of rightness' (FORs) and 'feelings of error' (FOEs).

The broader implications of this constructive work are threefold.

First, it supplements the taxonomy of extant theories of intuition to include what we might simply call 'the epistemic feelings theory of intuition'. While such a theory need not be taken as a direct rival to other accounts of intuitions – since, as previously discussed, intuition-talk likely refers to many different, diverse phenomena – it does provide researchers who focus on intuition with a new tool for their toolbox. For instance, it means that epistemic feelings are a new candidate to consider when endeavouring to interpret what the evidential basis of specific 'intuition-driven' philosophical arguments actually is; rather than just considering 'the usual suspects', such as beliefs or doxastic temptations, we can now also consider that it might be epistemic feelings. Armed with this will, in turn, hopefully lead to better insights among metaphilosophers about the evidential basis of specific philosophical arguments. At the very least, I have added e-feelings to the set of things that modern philosophers sometimes call 'intuitions', and rely upon evidentially, thereby helping to flesh-out our understanding of contemporary philosophical method.

Second, like the destructive portions of my research, the constructive portions have implications for how we should perform experimental investigations of intuitions. Narrowly, when trying to answer whether laypeople have the intuitions that contextualists claim they do, and/or whether said intuitions are sufficiently truth-tracking, we should now be trying to design experiments that test for the presence and/or reliability of epistemic feelings *specifically*, not, say, the beliefs or credences of experimental subjects. More widely, those who aspire to undertake such empirical investigations on other first-order philosophical claims should seriously consider the possibility that the relevant type of intuition is an epistemic feeling. For, while it's certainly possible that other philosophers, arguing in other areas, for different theories, do not, unlike these contextualists, appeal to e-feelings, the possibility that they are doing precisely that should at least be considered, now that we know that this is something that philosophers sometimes do.

Third, my new theory reshapes the broader landscape of the philosophy of intuition by reconciling 'seemings' accounts of intuition with 'minimalism'. To elaborate, on the one hand, minimalist accounts have tried to dispel the mystery surrounding intuitions by reducing them to another mental phenomenon. For example, as we've seen, 'doxasticism' is the position that intuitions are merely beliefs (Williamson, 2005), or at least a proper sub-set of beliefs, such as, specifically, spontaneous ones, common sense ones (Cappelen, 2012, p.71-75; Brogaard, 2013a; Parsons, 1995, p.59), or ones shared between the author and their, real or imagined, interlocutor(s)

(Cappelen, 2012). However, to reiterate, doxasticism is at odds with the non-committal nature of the intuitions in question: since it's possible to intuit *p* without believing *p*, the intuition that *p* can't just be a belief. In general, the problem with extant minimalist accounts – which was especially apparent in Chapter 3 – is that they postulate the *wrong* phenomenon as the mental state of intuition.

On the other hand, advocates of the 'intellectual seemings' theory of intuition claim that intuitions are a *sui generis* variety of 'perceptual seeming'. Unlike epistemic seemings, such as beliefs and confidences, which are committal and typically respond to changes in the subject's evidential corpus, perceptual seemings, such as visual experiences, are non-committal and recalcitrant. Intuitions, the theory goes, are perceptual seemings with their basis in reflection, rather than experience (Brogaard, 2013a). Brogaard (2013c), for example, says that intellectual seemings, or 'intuitions', have their source in what she calls 'armchair reflection' – which involves bringing to bear "both a priori principles and past experience" (p.9) – rather than sensory experiences. The standard intellectual seemings theories therefore respect the similarities between intuitions and experiences, but recognise the difference in their aetiology, and posit them as a unique, or *sui generis* entity.

However, while the intellectual seemings view acknowledges the aetiological differences between intuitions and experiences, it fails to say anything substantive about what this 'other' seeming actually is. Perceptual experiences are meant to be an example of the same, broad category of seeming, and intuitions are meant to be another, but we are left asking 'what is this *other* instance of a perceptual seeming, if it's not an experience?', and no answer is forthcoming. The intellectual seemings model has only identified some shared features of intuitions and experiences, and not told us what mental state is the former. In summation, while minimalist accounts posit the wrong state to be an intuition, *sui generis* accounts, like seemings theories, have a habit of not positing any state at all.

One advantage of my novel theory is that it strikes a happy medium between these minimalist and non-reductionist accounts of intuitions. On the one hand, e-feelings are a very plausible candidate for being the mental state of intuition, for the reasons given in this thesis. However, this view doesn't fall foul to the same counterexamples as other reductionist accounts, since, for example, e-feelings entail neither belief nor confidence in their content. On the other hand, rather than simply observing that intuitions are akin to perceptual experiences, but with a different aetiology – as extant seemings models are wont to do – we can now point to the specific phenomenon with that aetiology: e-feelings. Like standard seemings models, I've respected both the differences and similarities between intuitions and experiences, but, unlike standard seemings

models, identified the mental state of intuition in which those properties actually adhere, thereby providing a more substantive account.

By identifying intuitions with a certain variety of seemings, and seemings with a sub-category of an already-acknowledged type of mental state – epistemic feelings – I have essentially provided a minimalist or reductionist version of the *seemings* account of intuitions. Doing so circumvents the weaknesses currently found in both minimalist and *sui generis* theories. On the one hand, my minimalist account doesn't succumb to the same counterexamples that doom other reductionist theories. On the other hand, by pointing to a specific type of mental state, one that is already well-recognised in other fields of academic enquiry, I have not left my readers wondering what this mysterious sounding 'seeming' is meant to actually be.

One might also ask what implications my constructive findings have, not for the theory of intuition and experimental metaphilosophy more broadly, but for the plausibility of epistemic contextualism specifically. To that, I'll start by saying that to discern the actual consequences of my metaphilosophical theories for first-order epistemological views would take far more additional research than there is space for here, and that any such judgements made at this stage would likely be premature. That said, there are some potential implications that we can speculate about, and that provide promising avenues for future research.

For instance, as we've seen, contextualists treat the relevant intuitions as direct, first-order evidence for their content. However, this approach may be based on a misconception. Because e-feelings are an example of low-level (or experience-based), not high-level (or theory-based) metacognition, their content is entirely first-order (non-metarepresentational): when *S* thinks *p*, and finds *p* intuitive, the content of their FOR is simply *p*, and not anything like 'my thought that *p* was fluent', or 'my reasoning was sound'. In effect, this has camouflaged these intuitions for decades, masking their metacognitive nature. Although their content is first-order, they're feelings toward that first-order content triggered by higher-order facts about *S*' cognitive processes, not by the first-order 'truth makers' or 'justifiers' of *p*. Accordingly, intuitions, of this variety, at least, might be higher-order evidence instead.

To elaborate, while first-order evidence, *E*, is evidence that bears directly upon some hypothesis, *h*, higher-order evidence is evidence that bears directly upon *E* (Christensen, 2010; Kelly, 2010). For instance, that the weather forecast says it will be sunny tomorrow is evidence in favour of the proposition 'it will be sunny tomorrow'. By contrast, higher-order evidence is evidence that bears directly upon evidence. Higher-order evidence includes evidence for or against the existence of *E*, and evidence about how well *E* supports *h*. For instance, the fact that the weather forecast predicts sun tomorrow is evidence that forecasters have in their possession meteorological

evidence to that effect, and is thus evidence of the existence of such evidence. Alternatively, if *S* learns, say, that the forecast is often mistaken, this is evidence about how well said evidence supports the prediction. Plausibly, higher-order evidence also serves as indirect evidence of propositions. For example, evidence of meteorological evidence that it will be sunny tomorrow is also, indirectly, evidence that it will be sunny tomorrow.

Facts about *S*' own cognitive dispositions can also serve as higher-order evidence. *S* might, for instance, believe that it will be sunny tomorrow, but learn that they're prone to overestimating the chances of nice weather on the basis of what they read in the forecast. This self-knowledge is evidence about the quality of *S*' evidence; specifically, it suggests that whatever evidence they have doesn't support their prediction about the weather as strongly as they initially took it to. Again, this, in turn, is indirect evidence against their prediction: it suggests that the likelihood of it being sunny is lower than they initially thought.

One way to acquire higher-order evidence of the latter variety is, I argue, *via* e-feelings. To elaborate, because e-feelings are triggered by features of *S*' cognition, *S* can infer backwards, from said experiences, to those features. For example, biased reasoning is, again, especially prone to producing FORs, and thus, if *S* experiences a FOR when they judge that *h*, this suggests that said judgement may have been guided by a mental bias, rather than by an objective consideration of their evidential *corpus*. From this self-knowledge, *S* can infer that said *corpus* does not, perhaps, support *h* as strongly as it initially seemed to when they were judging *h*. Again, this higher-order evidence also serves as indirect evidence for or against *h* itself: *S*' self-knowledge suggests that *h* is less likely than they first thought. In other words, certain kinds of intuitions, rather than being first-order evidence for their content, could be higher-order evidence against their content: a radical inversion of how contextualists, at least, have been treating them.

More generally, because e-feelings are sensitive to facts about cognition, rather than to the truth-makers or justifiers of their content, they, at best, serve as indirect, higher-order evidence for, or against, that content. In other words, it might make more sense to use these particular intuitions to first learn facts about our *thinking*, and then infer facts about our *thoughts* from those. This discovery has the potential to revolutionise the role of said intuitions in philosophical method: philosophers will have to start treating their intellectual perceptual seemings as a very different kind of evidence, and assessing them along those lines (certainly, the methodological discussion more generally has, thus far, primarily concerned itself with the question of whether intuitions are good first-order evidence for propositions). Moreover, it has substantial consequences for, not just how these intuitions should be used in philosophy, but how we should understand their role in human rationality more generally.

So, my constructive findings suggest that a certain kind of intuition can be a source of self-knowledge, and then serve as higher-order evidence, and this has implications both for contextualists specifically, and for how this kind of intuition should be used and understood more widely, both within philosophy and without.

In summation, my research has several profound implications for the field of metaphilosophy in particular, as well as the methods of first-order philosophy more generally, and beyond. In the final section of this thesis, I'll briefly touch upon what future avenues of research these findings invite.

6.2 Future Research

In light of the above findings, and their implications, there are several avenues of future research now open to us. Perhaps most obviously, now that we have a much clearer picture of how intuitions get used in at least some seminal philosophical literature, and what those intuitions, at least, actually are, we have the beginnings of a more informed, detailed picture of modern philosophical method. Thus, future metaphilosophical research can now engage more confidently with that picture. That is, we now have some idea of what to measure, and what to assess, when trying to determine the answers to the more evaluative metaphilosophical research questions, namely:

C) Are the claims philosophers make about what people intuit accurate?
D) Are intuitions good evidence?

Indeed, we now have new, more specific ways of exploring the reliability of (some kinds of) philosophical intuitions. As we've seen throughout the thesis, there already exists substantial empirical data pertaining to the aetiology and reliability of e-feelings: what triggers them, how good they are at detecting conflict and bias, and so on. Thus, my work furnishes us with a new way to approach the empirical evaluation of philosophical method: we can utilise this data to assess the reliability of intellectual perceptual seemings in different contexts. Accordingly, our empirical investigation into the reliability of these philosophical intuitions, at least, can be expanded and refined.

Moreover, as I explained in the previous section, my findings suggest treating certain intuitions as higher-, not first-order evidence for their content. However, the exact upshots of construing intuitions as e-feelings will depend on many, intricate and overlapping factors, such as how we conceptualise higher-order evidence, and how higher-order evidence bears upon the normativity of first-order judgements. Future research could focus on these particulars, assessing the

implications of various theories of higher-order evidence and metacognition for philosophical method.

More generally, the research I've detailed in this thesis provides an approximate model for carrying out similarly cautious, narrowly focussed work on centrality in other areas of influential philosophical literature.

Bibliography

Alexander, J., Gonnerman, C. and Waterman, J. (2014) 'Salience and Epistemic Egocentricism: An Empirical Study' in Beebe, J. (ed.) *Advances in Experimental Epistemology*, Continuum

Alter, A. L., Oppenheimer, D. M., Epley, N. and Eyre, R. N. (2007) 'Overcoming Intuition: Metacognitive Difficulty Activates Analytic Reasoning' in *Journal of Experimental Psychology: General*, Vol. 136, pp. 569–576

Alter, A. L. and Oppenheimer, D. M. (2009) 'Uniting the Tribes of Fluency to Form a Metacognitive Nation' in *Personality and Social Psychology Review*, Vol. 13, pp. 219-235

Andow, J. (2015a) 'How 'Intuition' Exploded' in *Metaphilosophy*, Vol. 46, pp. 189–212

Andow, J. (2015b) 'Thin, Fine, and with Sensitivity: A Metamethodology of Intuitions' in *Review of Philosophy and Psychology* (online first), pp. 1–21

Andow, J. (2017) 'Intuition-talk: Virus or Virtue?' in *Philosophia*, Vol. 45, pp. 523-531

Arango-Muñoz, S. (2011) 'Two Levels of Metacognition' in *Philosophia*, Vol. 39, pp. 71-82

Arango-Muñoz, S. (2014) 'The Nature of Epistemic Feelings' in *Philosophical Psychology*, Vol. 27, pp. 193-211

Austin, J. L. (1946) 'Other Minds' in *Proceedings of the Aristotelian Society*, Supplementary Vol. 20, pp. 122-197

Balcomb, F. K. and Gerken, L. (2008) 'Three-Year-Old Children can Access their Own Memory to Guide Responses on a Visual Matching Task' in *Developmental Science*, Vol. 11, pp. 750-760

Baz, A. (2012) 'Must Philosophers Rely on Intuitions?' in *The Journal of Philosophy*, Vol. 109, pp. 316-337

Bealer, G. (1996) 'On the Possibility of Philosophical Knowledge' in *Philosophical Perspectives*, Vol. 10, pp. 1–34

Bealer, G. (1998) 'Intuition and the Autonomy of Philosophy' in DePaul, M. and Ramsey, W. (eds.) *Rethinking Intuition: The Psychology of Intuition and its Role in Philosophical Inquiry*, New York: Rowman and Littlefield, pp. 201–240

Bealer, G. (1999) 'A Theory of the A Priori' in Tomberlin, J. (ed.) *Philosophical Perspectives 13: Epistemology*, pp. 29–55

Bibliography

Bealer, G. (2000) 'A Theory of the a priori' in *Pacific Philosophical Quarterly*, Vol. 81, pp. 1-30

Bealer, G. (2008) 'Intuition and Modal Error' in Smith, Q. (ed.) *Epistemology: New Essays*, Oxford University Press

Bengson, J. (2014) 'How Philosophers Use Intuition and 'Intuition'' in *Philosophical Studies*, Vol. 171, pp. 555–76

Bengson, J. (2015) 'The Intellectual Given' in *Mind*, Vol. 124, pp. 707-760

Bermùdez, J. L. (2009) 'Mindreading in the Animal Kingdom' in Lurz, R. (ed.) *The Philosophy of Animal Minds*, Cambridge: Cambridge University Press, pp. 145-164

Blome-Tillmann, M. (2009) 'Knowledge and Presuppositions' in *Mind*, Vol. 118, pp. 241-294

Blome-Tillmann, M. (2013) 'Contextualism and the Knowledge Norms' in *Pacific Philosophical Quarterly*, Vol. 94, pp. 89-100

Boghossian, P. (2009) 'Virtuous Intuitions: Comments on Lecture 3 of Ernest Sosa's a Virtue Epistemology' in *Philosophical Studies*, Vol. 144, pp. 111-119

BonJour, L. (1985) *The Structure of Empirical Knowledge*, Harvard University Press

BonJour, L. (1998) *In Defense of Pure Reason*, Cambridge: Cambridge University Press

Bonner, C. and Newell, B. R. (2010) 'In Conflict with Ourselves? An Investigation of Heuristic and Analytic Processes in Decision Making' in *Memory and Cognition*, Vol. 38, pp. 186-196

Brogaard, B. (2012) 'Do We Perceive Natural Kind Properties?' in *Philosophical Studies*, Vol. 162, pp. 35-4

Brogaard, B. (2013a) 'It's Not What it Seems. A Semantic Account of 'Seems' and Seemings' in Cappelen, H. and Shapiro, S. (eds.) *Contextualism and Relativism (Special Issue of Inquiry)*

Brogaard, B. (2013b) 'Phenomenal Seemings and Sensible Dogmatism' in Tucker, C. (ed.) *Seemings and Justification*, Oxford: Oxford University Press

Brogaard, B. (2013c) 'Intuitions as Intellectual Seemings' in Sosa, D. (ed.) *Symposium on Herman Cappelen's Philosophy without Intuitions*, *Analytic Philosophy*, Vol. 55, pp. 382-393

Brown, R., and McNeill, D. (1966) 'The 'tip of the tongue' Phenomenon' in *Journal of Verbal Learning and Behavior*, Vol. 5, pp. 325-337

Brown, J. (2005) 'Comparing Contextualism and Invariantism on The Correctness of Contextualist Intuitions' in *Grazer Philosophische Studien*, Vol. 69, pp. 136-164

Buckwalter, W. (2010) 'Knowledge Isn't Closed on Saturday: A Study in Ordinary Language' in *The Review of Philosophy and Psychology*, Vol. 1, pp. 395-406

Buckwalter, W. (2012) 'Non-Traditional Factors in Judgments about Knowledge' in *Philosophy Compass*, Vol. 4, pp. 278-289

Buckwalter, W. and Schaffer, J. (2015) 'Knowledge, Stakes, and Mistakes' in *Noûs*, Vol. 49, pp. 201-234

Burge, T. (1979) 'Individualism and the Mental' in *Midwest Studies in Philosophy*, Vol. 4, pp. 73-122

Cappelen, H. (2012) *Philosophy Without Intuitions*, Oxford: Oxford University Press

Cappelen, H. (2014) 'Replies to Boghossian, Brogaard, and Richard' in *Symposium on Herman Cappelen's Philosophy without Intuitions*, *Analytic Philosophy*, Vol. 55, pp. 407-421

Cappelen, H. (2015) 'Replies to Weatherson, Chalmers, Weinberg, and Bengson' in *Symposium on Herman Cappelen's Philosophy without Intuitions, Philosophical Studies*, pp. 577-600

Carruthers, P. (2008) 'Meta-Cognition in Animals: A Skeptical Look' in *Mind and Language*, Vol. 23, pp. 58–89

Carruthers, P. (2009) 'How we Know our own Minds: The Relationship Between Mindreading and Metacognition' in *Behavioral and Brain Sciences*, Vol. 32, pp. 1-18

Carruthers, P. (2010) 'Introspection: Divided and Partly Eliminated' in *Philosophy and Phenomenological Research*, Vol. 80, pp. 76-111

Carruthers, P. and Ritchie, J. B. (2012) 'The Emergence of Metacognition: affect and uncertainty in animals' in Beran, M. J., Brandl, J. L., Perner, J. and Proust, J. (eds.) *Foundations of Metacognition*, Oxford: Oxford University Press, pp.76-93

Carruthers, P. (2017) 'Are Epistemic Emotions Metacognitive?' in *Philosophical Psychology*, Vol. 30, pp. 58-78

Carver, C. S. and Scheier, F. M. (1998) *On the Self-Regulation of Behavior*, Cambridge University Press

Carver, C. S. (2003) 'Pleasure as a Sign You Can Attend to Something Else: Placing Positive Feelings within a General Model of Affect' in *Cognition and Emotion*, Vol. 17, pp. 241-261

Chalmers, D. J. (1996) *The Conscious Mind: In Search of a Fundamental Theory*, New York and Oxford: Oxford University Press

Chalmers, D. J. (2014) 'Intuitions in Philosophy: A Minimal Defence' in *Philosophical Studies*, Vol. 171, pp. 535-544

Chisholm, R. (1957) *Perceiving: A Philosophical Study*, Ithaca: Cornell University Press

Chisholm, R. (1982) *Foundations of Knowing*

Christensen, D. (2004) *Putting Logic in its Place: Formal Constraints on Rational Belief*, Oxford: Oxford University Press

Christensen, D. (2010) 'Higher-Order Evidence' in *Philosophy and Phenomenological Research*, Vol. 81, pp. 185–215

Chudnoff (2011) 'What Intuitions Are Like' in *Philosophy and Phenomenological Research*, Vol. 82, pp. 625–654

Cohen, S. (1986) 'Knowledge and Context' in *The Journal of Philosophy*, Vol. 83, pp. 574-583

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

Cohen, S. (1998) 'Contextualist Solutions to Epistemological Problems: Skepticism, Gettier, and the Lottery' in *Australasian Journal of Philosophy*, Vol. 76, pp. 289-306

Cohen, S (1999) 'Contextualism, Skepticism, and The Structure of Reasons' in *Philosophical Perspectives, Vol. 13, Epistemology*, pp. 57–89

Cohen, S (2005) 'Knowledge, Speaker, Subject' in *The Philosophical Quarterly*, Vol. 55, pp. 199–212

Conee, E. (1998) 'Seeing the Truth' in *Philosophy and Phenomenological Research*, Vol. 58, pp. 847–57

Cohnitz, D. (2012) *Review of Philosophy Without Intuitions*

Costermans, J., Lories, G., and Ansay, C. (1992) 'Confidence Level and Feeling of Knowing in Question Answering: The Weight of Inferential Processes' in *Journal of Experiential Psychology: Learning, Memory, and Cognition*, Vol. 18, pp. 142–150

Crane, T. (1992) 'The Nonconceptual Content of Experience' in Crane, T. (ed.) *The Contents of Experience*, Cambridge: Cambridge University Press

Crane, T. (2001) *Elements of the Mind*, Oxford: Oxford University Press

Davidson, D. (1982) 'Rational Animals in *Dialectica*, Vol. 36, pp. 317-327

De Cruz, H. (2014) 'Where Philosophical Intuitions Come From' in *Australasian Journal of Philosophy*, Vol. 93, pp. 233-249

De Neys, W. (2006) 'Automatic-heuristic and Executive-analytic Processing During Reasoning: Chronometric and Dual-task Considerations' in *The Quarterly Journal of Experimental Psychology*, Vol. 59, pp. 1070–1100

De Neys, W. and Glumicic, T. (2008) 'Conflict Monitoring in Dual Process Theories of Thinking' in *Cognition*, Vol. 106, pp. 1248-1299

De Neys, W. (2012) 'Bias and Conflict: A Case for Logical Intuitions' in *Perspectives on Psychological Science*, Vol. 7, pp. 28–38

DeRose, K. (1992) 'Contextualism and Knowledge Attributions' in *Philosophy and Phenomenological Research*, Vol. 52, pp. 913-929

DeRose, K. (1995) 'Solving the Skeptical Problem' in *Philosophical Review*, Vol. 104, pp. 1-52

DeRose, K. (1999) 'Contextualism: An Explanation and Defense' in Greco, J. and Sosa, E. (eds.) *The Blackwell Guide to Epistemology*, Blackwell Publishers

DeRose, K. (2000) 'Now You Know It, Now You Don't' in *Proceedings of the Twentieth World Congress of Philosophy*, Vol. V, Epistemology, pp. 91–106

DeRose, K. (2009) *The Case for Contextualism: Knowledge, Skepticism, and Context, Vol. 1*, Oxford: Oxford University Press

DeRose, K. (2017) *The Appearance of Ignorance: Knowledge, Skepticism, and Context, Vol. 2*, Oxford: Oxford University Press

de Sousa, R. (1987) *The Rationality of Emotion*, Cambridge, MA: MIT Press

de Sousa, R. (2008) 'Epistemic Feelings' in Brun, G., Doguoglu, U., and Kuenzle, D. (eds.) *Epistemology and Emotions*, Hampshire, UK: Ashgate pp. 185-204

Deutsch, M. (2010) 'Intuitions, Counter-Examples, and Experimental Philosophy' in *Review of Philosophy and Psychology*, Vol. 1, pp. 447-460

Deutsch, M. (2015) *The Myth of the Intuitive: Experimental Philosophy and Philosophical Method*

Döring, S. (2003) 'Explaining Action by Emotion' in *Philosophical Quarterly*, Vol. 53, pp. 214-230

Dretske, F. (1995) *Naturalizing the Mind*, Cambridge, MA: MIT Press

Dretske, F. (2005) 'The Case Against Closure' in Steup, M. and Sosa, E. (eds.) *Contemporary Debates in Epistemology*, pp. 1–26

Dunlosky, J., and Metcalfe, J. (2009). *Metacognition*. Thousand Oaks, CA: Sage

Earlenbaugh, J. and Molyneux, B. (2009) 'Intuitions are inclinations to believe' in *Philosophical Studies*, Vol. 145, pp. 89–109

Efklides, A. (2006) 'Metacognition and Affect: What Can Metacognitive Experiences tell us About the Learning Process?' in *Educational Research Review*, Vol. 1, pp. 3-14

Evans, J. St. B. T. (2003) 'In Two Minds: Dual-process Accounts of Reasoning' in *Trends in Cognitive Sciences*, Vol. 7, 454-459

Evans, J. St. B. T and Curtis-Holmes, J. (2005) 'Rapid Responding Increases Belief Bias: Evidence for the Dual-Process Theory of Reasoning' in *Thinking and Reasoning,* Vol. 11, pp. 382–389

Evans, J. St. B. T. (2008) 'Dual-Processing Accounts of Reasoning, Judgment and Social Cognition' in *Annual Review of Psychology*, Vol. 59, pp. 255-278

Evans, J. St. B. T. (2010) 'Intuition and Reasoning: a Dual-Process Perspective' in *Psychological Inquiry*, Vol. 21, pp. 313–326

Evans, J. St. B. T. and Stanovich, K. E. (2013) 'Dual-Process Theories of Higher Cognition: Advancing the Debate' in *Perspectives on Psychological Science*, Vol. 8, pp. 223–241

Feltz, A. and Zarpentine, C. (2010) 'Do You Know More When it Matters Less?' in *Philosophical Psychology*, Vol. 23, pp. 683-706

Finucane, M. L., Alhakami, A., Slovic, P., and Johnson, S. M. (2000) 'The Affect Heuristic in Judgments of Risks and Benefits' in *Journal of Behavioral Decision Making*, Vol. 13, pp. 1-17

Fischer, E. (2014) 'Philosophical Intuitions, Heuristics, and Metaphors' in *Synthese*, Vol. 191, pp. 569-606

Fischer, E. and Collins, J. (2015) 'Rationalism and Naturalism in the Age of Experimental Philosophy' in Fischer, E. and Collins, J. (eds.) *Experimental Philosophy, Rationalism, and Naturalism. Rethinking Philosophical Method*, Routledge, pp. 3-33

Fischer, E., Engelhardt, P. E. and Herbelot, A. (2015) 'Intuitions and Illusions: From Explanation and Experiment to Assessment' in Fischer, E. and Collins, J. (eds.) *Experimental Philosophy, Rationalism, and Naturalism. Rethinking Philosophical Method*, Routledge, pp. 259-292

Fischer, E. (2018) 'Review of Max Deutsch The Myth of the Intuitive: Experimental Philosophy and Philosophical Method' in *The Philosophical Review*, Vol. 127, pp. 413-418

Flavell, J. (1979) 'Metacognition and Cognitive Monitoring: A new area of Cognitive Developmental Inquiry' in *American Psychologist*, Vol. 34, pp. 906-911

Flavell, J. H. (1998) 'The Mind has a Mind of its Own: Developing Knowledge about Mental Uncontrollability' in *Cognitive Development*, Vol. 13, pp. 127-138

Foot, P. (1967) 'The Problem of Abortion and the Doctrine of the Double Effect' in *The Oxford Review*, Vol. 5, pp. 5-15

Gangemi, A., Bourgeois-Gironde, S. and Mancini, F. (2015) 'Feelings of Error in Reasoning – in Search of a Phenomenon' in *Thinking and Reasoning*, Vol. 21, pp. 383-396

Gettier, E. (1963) 'Is Justified True Belief Knowledge?' in *Analysis*, Vol. 23, pp. 121–123

Gilbert, D. T., Krull, D. S., and Malone, P. S. (1990) 'Unbelieving the Unbelievable: Some Problems in the Rejection of False Information' in *Journal of Personality and Social Psychology*, Vol. 59, pp. 601-613

Gilbert, D. T., Tafarodi, R.W., and Malone, P. S. (1993) 'You Can't Not Believe Everything You Read' in *Journal of Personality and Social Psychology*, Vol. 65, pp. 221-233

Gladwell, M. (2007) *Blink: The Power of Thinking Without Thinking*, New York, Black Bay Books

Glucksberg, S., and McCloskey, M. (1981) 'Decisions About Ignorance: Knowing that you don't Know' in *Journal of Experimental Psychology: Human Learning and Memory*, Vol. 7, pp. 311-325

Goldie, P. (2000) *The Emotions: A Philosophical Exploration*, Oxford: Clarendon

Goldie, P. (2002) 'Emotion, Feelings and Intentionality' in *Phenomenology and the Cognitive Sciences*, Vol. 1, pp. 235-254

Goldman, A.I. (1986) *Epistemology and Cognition*, Harvard University Press

Goldman, A.I. (2007) ‘Philosophical Intuitions: Their Target, Their Source, and Their Epistemic Status’ in *Grazer Philosophische Studien*, Vol. 74, pp. 1–26

Goodman, N. (1965) *Fact, Fiction and Forecast*

Gopnik, A. and E. Schwitzgebel (1998) ‘Whose Concepts are They, Anyway? The Role of Philosophical Intuition in Empirical Psychology’ in DePaul, M. R. and Ramsey, W. (eds.) *Rethinking Intuition*, Lanham, MD: Rowman and Littlefield, pp. 75-91

Hambourger, R. (1987) ‘Justified Assertion and the Relativity of Knowledge’ in *Philosophical Studies*, Vol. 51, pp. 241-269

Hampton, R. (2001) ‘Rhesus Monkeys Know when they Remember’ in *Proceedings of the National Academy of Sciences*, Vol. 98, pp. 5359-5362

Handley, S. J., Newstead, S. E., and Trippas, D. (2010) ‘Logic, Beliefs, and Instruction: A test of the Default Interventionist Account of Belief Bias’ in *Learning, Memory, and Cognition*, Vol. 37, pp. 28–43

Harman, G. (1968) ‘Knowledge, Inference, and Explanation’ in *American Philosophical Quarterly*, Vol. 5, pp. 164-173

Hawthorne, J. (2004) *Knowledge and Lotteries*, Oxford: Oxford University Press

Huemer, M. (2001) ‘Skepticism and the Veil of Perception’

Huemer, M. (2007) ‘Compassionate Phenomenal Conservatism’ in *Philosophy and Phenomenological Research*, Vol. 74, pp. 30-55

Huemer, M. (2008) ‘Revisionary Intuitionism’ in *Social Philosophy and Policy*, Vol. 25, pp. 368-392

Hofweber, T. (1999) ‘Contextualism and the Meaning-Intention Problem’ in *Cognition, Agency and Rationality*, Korta, K., Sosa, E. and Arrazola, X. (eds.), London: Kluwer, pp. 93-104

Ichikawa, J. and Jarvis, B. W. (2013) *The Rules of Thought*, Oxford: Oxford University Press

Ichikawa, J. (2014) ‘Who Needs Intuitions? Two Experimentalist Critiques’ in Booth, A. R. and Rowbottom, D. P. (eds.) *Intuitions*, Oxford University Press, pp. 232-256

Jackson, F. (1977) *Perception: A Representative Theory*, Cambridge: Cambridge University Press

Jameson, A., Nelson, T. O., Leonesio, R. J., and Narens, L. (1993) ‘The Feeling of Another Person’s Knowing’ in *Journal of Memory and Language*, Vol. 32, pp. 320-335

Kahneman, D. (2003) 'A Perspective on Judgment and Choice: Mapping Bounded Rationality' in *American Psychologist*, Vol. 58, pp. 697–720

Kahneman, D. and Frederick, S. (2005) 'A Model of Heuristic Judgment' in Holyoak, K. J. Morrison, R. (eds.) *The Cambridge Handbook of Thinking and Reasoning*, Cambridge: Cambridge University Press, pp. 67–293

Kahneman, D. (2011) *Thinking, Fast and Slow*

Jeffrey, R. (1965) *The Logic of Decision*, Chicago: Chicago University Press

Kagan, S. (1989) *The Limits of Morality*, Oxford University Press

Katz, J. (1981) *Language and Other Abstract Objects*

Kelley, C. M. and Lindsay, D. S. (1993) 'Remembering Mistaken for Knowing: Ease of Retrieval as a Basis for Confidence in Answers to General Knowledge Questions' in *Journal of Memory and Language*, Vol. 32, pp. 1–24

Kelly, T. (2010) 'Peer Disagreement and Higher Order Evidence' in Feldman and Warfield (eds.) *Disagreement*, Oxford: Oxford University Press, pp. 111–174

Koksvik, O. (2012) *Intuition*, Ph.D. Thesis, Australian National University, Available from: http://www.koksvik.net/files/koksvik_intuition_monochrome.pdf

Koriat, A. (1993) 'How do we Know that we Know? The Accessibility Model of the Feeling of Knowing' in *Psychological Review*, Vol. 100, pp. 609-639

Koriat, A. (1995) 'Dissociating Knowing and the Feeling of Knowing: Further Evidence for the Accessibility Model' in *Journal of Experimental Psychology: General*, Vol. 124, pp. 311-333

Koriat, A. and Levy-Sadot (1999) 'Processes Underlying Metacognitive Judgements: Information-based and Experience-based Monitoring of One's Own Knowledge' in Chaiken, S. and Trope, Y. (eds.) *Dual Process Theories in Social Psychology*, New York: Guilford, pp. 483-502

Koriat, A. (2000) 'The Feeling of Knowing: Some Metatheoretical Implications for Consciousness and Control' in *Consciousness and Cognition*, Vol. 9, pp. 149-171

Koriat, A., Bjork, R. A., Sheffer, L. and Bar, S. K. (2004) 'Predicting One's Own Forgetting: The Role of Experience-Based and Theory-Based Processes' in *Journal of Experimental Psychology: General*, Vol. 133, pp. 643-656

Koriat, A., Ma'ayan, H. and Nussinson, R. (2006) 'The Intricate Relationships Between Monitoring and Control in Metacognition: Lessons for the Cause-and-effect Relation Between Subjective Experience and Behavior' in *Journal of Experimental Psychology: General*, Vol. 135, pp. 36-69

Kornblith, H. (1998) 'The Role of Intuitions in Philosophical Inquiry' in DePaul, M. and Ramsay, W. (eds.) *Rethinking Intuition*, Lanham, MD: Rowman and Littlefield

Kripke, S. (1980) *Naming and Necessity*, Oxford: Blackwell Publishing

Kripke, S. A. (2011) *Philosophical Troubles: Collected Papers*, Vol. 1, Oxford University Press

Lehrer, K. (2000) *Theory of Knowledge*, 2nd Edition, Boulder: Westview Press

Lewis, D. (1979) 'Scorekeeping in a Language Game' in *Journal of Philosophical Logic*, Vol. 8, pp. 339-359

Lewis, D. (1983) *Philosophical Papers: Volume I*, New York: Oxford University Press

Lewis, D. (1996) 'Elusive Knowledge' in *Australasian Journal of Philosophy*, Vol. 74, pp. 549-567

Liao, M. (2008) 'A Defense of Intuitions' in *Philosophical Studies*, Vol. 140, pp. 247-262

Ludwig, K. (2007) 'The Epistemology of Thought Experiments: First Person Versus Third Person Approaches' in *Midwest Studies in Philosophy: Philosophy and the Empirical*, Vol. 31, pp. 128–159

Ludwig, K. (2010) 'Intuitions and Relativity' in *Philosophical Psychology*, Vol. 23, pp. 427-445

Lycan, W. (1998) 'Moral Facts and Moral Knowledge' in Lycan, W. (ed.) *Judgement and Justification*, Cambridge: Cambridge University Press

Lynch, M. (2006) 'Trusting Intuition' in Greenough, P. and Lynch, M. (eds.) *Truth and Realism*, Oxford: Oxford University Press

Milner, D. and Goodale, M. (1995) *The Visual Brain in Action*, New York: Oxford University Press

Moore, G. E. (1959) 'Proof of an External World', in *Philosophical Papers,* London: George Allen and Unwin, pp. 127-150

Nado, J. (2015) 'Intuition Deniers' in *Philosophical Studies*, Vol. 173, pp. 781-800

Nagel, J. (2011) 'The Psychological Basis of the Harman-Vogel Paradox' in *Philosophers' Imprint*, Vol. 11, pp. 1-28

Nagel, J. (2014) 'The Meanings of Metacognition' in *Philosophy and Phenomenological Research*, Vol. 89, pp. 710-718

Nelson, T. O. and Narens, L. (1980) 'Norms of 300 General-information Questions: Accuracy of Recall, Latency of Recall, and Feeling-of-Knowledge Ratings' in *Journal of Verbal Learning and Verbal Behavior*, Vol. 19, pp. 338-368

Nelson, T. and Narens, L. (1990) 'Metamemory: A Theoretical Framework and new Findings' in Bower, G. (ed.) *The Psychology of Learning and Information*, Vol. 26, pp. 125-174

Nichols, S., Stich, S. and Weinberg, J. (2003) 'Meta-skepticism: Meditations on Ethno-epistemology' in Luper, S. (ed.) *The Skeptics*, Aldershot, U.K.: Ashgate Publishing, pp. 227-247

Nussbaum, M. (2001) *Upheavals of Thought: The Intelligence of Emotions*, Cambridge: Cambridge University Press

Nussinson, R. and Koriat, A. (2008) 'Correcting Experience-Based Judgments: the Perseverance of Subjective Experience in the Face of the Correction of Judgment' in *Metacognition Learning*, Vol. 3, pp. 159-174

Panksepp, J. (1998) *Affective Neuroscience*, New York: Oxford University Press

Parsons, C. (1995) 'Platonism and Mathematical Intuition in Kurt Gödel's Thought' in *Bulletin of Symbolic Logic*, Vol. 1, pp. 44–74

Peacocke, C. (1983) Sense and Content: *Experience, Thought, and Their Relations*, Oxford: Oxford University Press

Peacocke, C. (2004) *The Realm of Reason*, Oxford: Oxford University Press

Peacocke, C. (2007) 'Mental Action and Self-awareness (I)' in McLaughlin, B. and Cohen, J. (eds.) *Contemporary Debates in the Philosophy of Mind*, Malden, MA: Blackwell Publishing, pp. 358-376

Peacocke, C. (2009) 'Mental Action and Self-awareness (II)' in O'Brien, L. and Soteriou, M. (eds.) *Mental Actions and Agency*, Oxford: Oxford University Press, pp. 192-214

Pennycook, G., Trippas, D., Handley, S. J., and Thompson, V. A. (2013) 'Base Rates: Both Neglected and Intuitive' in *Journal of Experimental Psychology: Learning, Memory, and Cognition*, Vol. 40, pp. 544-554

Phelan, M. (2013) 'Evidence that Stakes Don't Matter for Evidence' in *Philosophical Psychology*, Vol. 4, pp. 1-25

Piattelli-Palmarini, M. (1994) *Inevitable Illusions*, New York: Wiley

Pinillos, N. Á. (2012) 'Knowledge, Experiments and Practical Interests' in Brown, J. and Gerken, M. (eds.) *New Essays On Knowledge Ascriptions*, Oxford University Press

Plantinga, A. (1993) *Warrant and Proper Function*, Oxford University Press

Pollock, J. (1974) *Knowledge and Justification*, Princeton University Press

Price, R. (1758) 'A Review of the Principle Questions in Morals' in Raphael, D. D. (ed.) (1969) *The British Moralists 1650–1800, II*, Oxford: Clarendon Press, pp. 131–198

Prince, E., Bosk, C., and Frader, J. (1982) 'On Hedging in Physician–Physician Discourse' in di Pietro, J. (ed.) *Linguistics and the Professions*, pp. 83–97

Prinz, J. (2004) *Gut Reactions: A Perceptual Theory of Emotion*, Oxford: Oxford University Press

Proust, J. (2007) 'Metacognition and Metarepresentation: Is a Self-directed Theory of Mind a Precondition for Metacognition?' in *Synthese*, Vol. 159, pp. 271-295

Proust, J. (2009a) 'It There a Sense of Agency of Thought?' in O'Brien, L. and Soteriou, M. (eds.) *Mental Actions and Agency*, Oxford: Oxford University Press, pp. 253-279

Proust, J. (2009b) 'The Representational Basis of Brute Metacognition: A Proposal' in Lurz, R. (ed.) *The Philosophy of Animal Minds*, Cambridge: Cambridge University Press, pp. 165-183

Proust, J. (2012) 'Metacognition and Mindreading: One or Two Functions?' in M. Beran, Brandl, J., Perner, J., and Proust, J. (eds.) *Foundations of Metacognition*, New York: Oxford University Press, pp. 234-251

Proust, J. (2014) *The Philosophy of Metacognition: Mental Agency and Self-Awareness*, New York: Oxford University Press

Prowse Turner, J. A., and Thompson, V. (2009) 'The Role of Training, Alternative Models, and Logical Necessity in Determining Confidence in Syllogistic Reasoning' in *Thinking and Reasoning*, Vol. 15, pp. 69-100

Pryor, J. (2005) 'There is immediate justification' in Steup, M. and Sosa, E. (eds.) *Contemporary Debates in Epistemology*, John Wiley and Sons

Pust, J. (2000) *Intuitions as Evidence*, New York: Routledge

Pust, J. (2017) 'Intuition' in Zalta, E. N. (ed.) *Stanford Encyclopedia of Philosophy*, Available From: [https://plato.stanford.edu/archives/sum2017/entries/intuition]

Putnam, H. (1975) 'The Meaning of 'Meaning'' in Gunderson, K. (ed.) *Language, Mind and Knowledge, Minnesota Studies in the Philosophy of Science 7*, pp. 131-193

Putnam, H. (1981) 'Brains in a Vat' in *Reason, Truth, and History*

Ramsey, F. (1931) 'Truth and Probability' in Braithwhite, R. (ed.) *The Foundations of Mathematics and other Logical Essays*, London: Routledge and Kegan Paul, pp. 156-198

Rawls, J. (1971) *A Theory of Justice*, Cambridge, MA: Harvard University Press

Reder, L. M. and Ritter, F. E. (1992) 'What Determines Initial Feeling of Knowing? Familiarity with Question Terms, not with the Answer' in *Journal of Experimental Psychology: Learning, Memory, and Cognition*, Vol. 18, pp. 435-451

Reder, L. M. (ed.) (1996) *Implicit Memory and Metacognition*

Reid, T. (1785), in Brookes, D. R. (ed.) (2002) *Essays on the Intellectual Powers of Man,* Edinburgh: Edinburgh University Press

Richard, M. (2014) 'Analysis, Concepts, and Intuitions' in *Analytic Philosophy*, Vol. 55, pp. 394-406

Richard, M. (2019) *Meanings as Species*, Oxford: Oxford University Press

Rini, R. A. (2015) 'How Not to Test for Philosophical Expertise' in *Synthese*, Vol. 192, pp. 431-452

Roberts, M. J., and Newton, E. J. (2001) 'Inspection Times, the Change Task, and the Rapid-Response Selection Task' in *The Quarterly Journal of Experimental Psychology*, Vol. 54A, pp. 1031-1048

Roberts, R.C. (2009) 'The Sophistication of Non-human Emotion' in Lurz, R. (ed.) *The Philosophy of Animal Minds*, Cambridge: Cambridge University Press, pp. 218-236

Robinson, M. D., Johnson, J. T. and Herndon, F. (1997) 'Reaction Time and Assessments of Cognitive Effort as Predictors of Eyewitness Memory Accuracy and Confidence' in *Journal of Applied Psychology*, Vol. 82, pp. 416–425

Schiffer, S. (1996) 'Contextualist Solutions to Skepticism' in *Proceedings of the Aristotelian Society*, Vol. 96, pp. 317–333

Schiffrin, R. and Schneider, W. (1977) 'Controlled and Automatic Human Information Processing: I Detection, Search and Attention' in *Psychological Review*, Vol. 84, pp. 1–66

Schunn, C. D., Reder, L. M., Nhouyvanisvong, A., Richards, D. R. and Stroffolino, P. J. (1997) 'To Calculate or not to Calculate: A Source Activation Confusion Model of Problem Familiarity's Role in Strategy Selection' in *Journal of Experimental Psychology: Learning, Memory, and Cognition*, Vol. 23, pp. 3-29

Schwarz, N., Bless H., Strack, F., Klumpp, G., Rittenauer-Schatka, H., and Simons, A. (1991) 'Ease of Retrieval as Information: Another Look at the Availability Heuristic' in *Journal of Personality and Social Psychology*, Vol. 61, pp. 195-202

Schwartz, B. (2002) *Tip-of-the-Tongue States: Phenomenology, Mechanisms and Lexical Retrieval*, Mahwah, NJ: Lawrence Erlbaum Associates

Schwartz, B. and Metcalfe, J. (2010) 'Tip-of-the-Tongue (TOT) States: Retrieval, Behavior, and Experience' in *Memory and Cognition*, Vol. 39, pp. 737-749

Seager, W. (2002) 'Emotional Introspection' in *Consciousness and Cognition*, Vol. 11, pp. 666-687.

Searle, J.R. (1983) *Intentionality: An Essay in the Philosophy of Mind*, Cambridge: Cambridge University Press

Shafer, K. (2013) 'Perception and the Rational Force of Desire' in *The Journal of Philosophy*, Vol. 110 pp. 258-281

Siegel, S. (2010) 'Do Visual Experiences have Contents?' in Nanay, B. (ed.) *Perceiving the World*, Oxford: Oxford University Press

Simmons, J. P. and Nelson, L. D. (2006) 'Intuitive Confidence: Choosing Between Intuitive and Nonintuitive Alternatives' in *Journal of Experimental Psychology: General*, Vol. 135, pp. 409–428

Skyrms, B. (1980) 'Higher Order Degrees of Belief' in Mellor, D. H. (ed.) *Prospects for Pragmatism*, Cambridge: Cambridge University Press, pp. 109-137

Sloman, S. A. (1996) 'The Empirical Case for Two Systems of Reasoning' in *Psychological Bulletin*, Vol. 119, pp. 3–22

Sloman, S. A. (2002) 'Two Systems of Reasoning' in Griffin, G. D. and Kahneman, D. (eds.) *Heuristics and Biases: The Psychology of Intuitive Judgment*, New York, NY: Cambridge University Press, pp. 379–396

Smith, J. D. (2009) 'The Study of Animal Metacognition' in *Trends in Cognitive Sciences*, Vol. 13, pp. 389-396

Smith, T. (2000) *Viable Values*, Rowman and Littlefield

Solomon, R. (1993) *The Passions: Emotions and the Meaning of Life*, Indianapolis: Hackett

Solomon, R. (2004) *Thinking About Feeling*, Oxford: Oxford University Press

Sosa, E. (1998) 'Minimal Intuition' in DePaul, M. and Ramsey, W (eds.) *Rethinking Intuition: The Psychology of Intuition and Its Role in Philosophical Inquiry*, Rowman and Littlefield, pp. 257–269

Sosa, E. (2000) 'Skepticism and Contextualism' in *Philosophical Issues*, Vol. 10, pp. 1-18

Sosa, E. (2006) 'Intuitions and Truth' in Greenough, P. and Lynch, M. (eds.) *Truth and Realism*, Oxford: Oxford University Press

Sosa, E. (2007a) *A Virtue Epistemology: Apt Belief and Reflective Knowledge*, Oxford, Oxford University Press

Sosa, E. (2007b) 'Intuitions: Their Nature and Epistemic Efficacy' in *Grazer Philosophische Studien*, Vol. 74, pp. 51-67

Sosa, E. (2007c) 'Experimental Philosophy and Philosophical Intuition', in *Philosophical Studies*, Vol. 132, pp. 99-107

Sosa, E. (2009) 'Replies to Commentators on A Virtue Epistemology', *Philosophical Studies*, Vol. 144, pp. 137-47

Sosa, E. (2015) *Judgment and Agency*, Oxford: Oxford University Press

Srinivasan, A. (2015) 'Are We Luminous?' in *Philosophy and Phenomenological Research*, Vol. 90, pp. 294-319

Sripada, C. S. and Stanley, J. (2012) 'Empirical Tests of Interest-Relative Invariantism' in *Episteme*, pp. 3-26

Stanley, J. (2004) 'On the Linguistic Basis for Contextualism' in *Philosophical Studies*, Vol. 119, pp. 119-146

Stanley, J. (2007) *Knowledge and Practical Interests*, Oxford University Press

Stanovich, K. (1999) *Who is Rational?: Studies in Individual Differences in Reasoning*, Mahwah, NJ: Lawrence Erlbaum

Stanovich, K. E. and West, R. F. (2000) 'Advancing the Rationality Debate' in *Behavioral and Brain Sciences*, pp. 701-717

Stanovich, K. E. (2011) *Rationality and The Reflective Mind*, New York: Oxford University Press

Stitch, S. (2013) 'Do Different Groups Have Different Epistemic Intuitions? A Reply to Jennifer Nagel' in *Philosophy and Phenomenological Research*, Vol. 87, pp. 151-178

Strack, F., and Deutsch, R. (2004) 'Reflective and Impulsive Determinants of Social Behavior' in *Personality and Social Psychology Review*, Vol. 8, pp. 220–247

Tenenbaum, S. (2013) 'Guise of the Good' in LaFollette, H. (ed.) *The International Encyclopedia of Ethics*, Blackwell Publishing

Thomson, J. J. (1971) 'A Defense of Abortion' in *Philosophy and Public Affairs*, Vol. 1, pp. 47-66

Thompson, V. A. (2009) 'Dual-Process Theories: A Metacognitive Perspective' in Evans, J. and Frankish, K. (eds.) *In Two Minds: Dual Processes and Beyond*, Oxford: Oxford University Press, pp. 171-195

Thompson, V. A., Prowse Turner, J. A. and Pennycook, G. (2011) 'Intuition, Reason, and Metacognition' in *Cognitive Psychology*, Vol. 63, pp. 107–140

Thompson, V. A. and Morsanyi, K. (2012) 'Analytic Thinking: Do you Feel Like it?' in *Mind and Society*, Vol. 11, pp. 93-105

Thompson, V. A., Evans, J. St. B. T. and Campbell, J. I. C. (2013) 'Matching Bias on the Selection Task: It's Fast and it Feels Good' in *Thinking and Reasoning*, Vol. 19, pp. 431-452

Thompson, V. A. and Johnson, S. J. (2014) 'Conflict, Metacognition, and Analytic Thinking' in *Thinking and Reasoning*, Vol. 20, pp. 215-244

Tsujii, T. and Watanabe, S. (2010) 'Neural Correlates of Belief-Bias Reasoning Under Time-Pressure: A Near-Infrared Spectroscopy Study' in *NeuroImage*, Vol. 50, pp. 1320–1326

Tye, M. (1995) *Ten Problems about Consciousness: A Representational Theory of the Phenomenal Mind*, Cambridge, MA: MIT Press

Tye, M. (2008) 'The Experience of Emotion: An Intentionalist Theory' in *Revue Internationale de Philosophi*e, Vol. 243, pp. 25-50

Tye, M. (2009) 'Representationalist Theories of Consciousness' in McLaughlin, B. and Beckermann, A. (eds.), *The Oxford Handbook of Philosophy of Mind*. Oxford: Oxford University Press

Unger, P. (1975) *Ignorance: A Case for Scepticism*, Oxford: Oxford University Press

Unkelbach, C. (2007) 'Reversing the Truth Effect: Learning the Interpretation of Processing Fluency in Judgments of Truth' in *Journal of Experimental Psychology*: *Learning, Memory and Cognition*, Vol. 33, pp. 219-230

Van Inwagen, P. (1997) 'Materialism and the Psychological-Continuity Account of Personal Identity' in *Philosophical Perspectives*, Vol. 11, pp. 305–319

Van Roojen, M. (2016) 'Moral Cognitivism vs. Non-Cognitivism' in Zalta, E. N. (ed.) *The Stanford Encyclopedia of Philosophy* (Winter 2016 Edition), Available from: [https://plato.stanford.edu/archives/win2016/entries/moral-cognitivism/]

Vernon, D. and Usher, M. (2003) 'Dynamics of Metacognitive Judgments: Pre-and Post Retrieval Mechanisms' in *Journal of Experimental Psychology: Learning, Memory, and Cognition*, Vol. 29, pp. 339-346

Volz, K. G., Rübsamen, R. and Von Cramon, D. Y. (2008) 'Cortical Regions Activated by the Subjective Sense of Perceptual Coherence of Environmental Sounds: A Proposal for a Neuroscience of Intuition' in *Cognitive, Affective, and Behavioral Neuroscience*, Vol. 8, pp. 318-328

Walsh, M. M. and Anderson, J. R. (2009) 'The Strategic Nature of Changing Your Mind' in *Cognitive Psychology*, Vol. 58, pp. 416-440

Wason, P. C. (1959) 'The Processing of Positive and Negative Information' in *Quarterly Journal of Experimental Psychology*, Vol 11, pp. 92-107

Weatherson, B. (2013) 'Centrality and Marginalisation' in Sosa, D. (ed.) *Symposium on Herman Cappelen's Philosophy Without Intuitions*, Analytic Philosophy

Weinberg, J. M., Nichols, S. and Stitch, S. (2001) 'Normativity and Epistemic Intuitions' in *Philosophical Topics*, Vol. 29, pp. 429-460, reprinted in Knobe, J. and Nichols, S (eds.) (2008) *Experimental Philosophy*, New York: Oxford University Press, pp. 17-36

Whittlesea, B. (1993) 'Illusion of Familiarity' in *Journal of Experimental Psychology*, Vol. 19, pp. 1235-1253

Whittlesea, B. and Williams, L. (2001) 'Source of the Feeling of Familiarity: The Discrepancy-Attribution Hypothesis' in *Journal of Experimental Psychology*, Vol. 26, pp. 547-565

Williams, B. (1970) 'The Self and the Future' in *Philosophical Review*, Vol. 79, pp. 161-180

Williamson, T. (2000) *Knowledge and its Limits*, New York: Oxford University Press

Williamson, T. (2005) 'Contextualism, Subject-Sensitive Invariantism and Knowledge of Knowledge' in *The Philosophical Quarterly*, Vol. 55, pp. 213-235

Williamson, T. (2007) *The Philosophy of Philosophy*, Oxford: Blackwell

Williamson, T. (2011) 'Philosophical Expertise and the Burden of Proof' in *Metaphilosophy*, Vol. 42, pp. 215–229

Yeung, H. H., Denison, S. and Johnson, S. P. (2016) 'Infants Looking to Surprising Events: When Eye-Tracking Reveals More than Looking Time' in *Plos One*

Zajonc, R. B. and Rajecki, D. W. (1969) 'Exposure and Affect: A Field Experiment' in *Psychonomic Science*, Vol. 17, pp. 216-217