



Minerva Access is the Institutional Repository of The University of Melbourne

Author/s:

Tuckwell, William

Title:

Non-ideal Epistemic Contextualism

Date:

2021

Persistent Link:

<https://hdl.handle.net/11343/293058>

Terms and Conditions:

Terms and Conditions: Copyright in works deposited in Minerva Access is retained by the copyright owner. The work may not be altered without permission from the copyright owner. Readers may only download, print and save electronic copies of whole works for their own personal non-commercial use. Any use that exceeds these limits requires permission from the copyright owner. Attribution is essential when quoting or paraphrasing from these works.

Non-ideal Epistemic Contextualism

William Tuckwell

ORCID: 0000-0002-1815-4871

Doctor of Philosophy

July 2021

School of Historical and Philosophical Studies

A thesis submitted in total fulfilment of the requirements of the degree of Doctor of
Philosophy.

Abstract

Epistemic contextualists claim that in order for knowledge ascribing sentences, i.e., sentences of the form ‘S knows that p’, to be true S must meet different epistemic standards in different contexts. Some contextualists, those who I’ll label conversational contextualists, claim that speakers can change which standards are contextually operative by making certain conversational moves, e.g., by raising an alternative to the proposition that features in the knowledge ascription to salience. The examples that these theorists work with and the theories that they develop off the back of them tend to build in the assumption that all speakers have equal abilities to pull off standards raising moves. But this is an idealisation that neglects the fact that in our non-ideal world, speakers have unequal abilities to successfully pull off conversational moves, where these unequal abilities vary with unfairly and unjustly distributed power. The project of this thesis is to show that paying greater attention to this fact has three substantial payoffs.

Firstly, it allows us to see unrecognised problems with various existing forms of contextualism. In chapter 3 I argue that Keith DeRose’s gap view is vulnerable to an objectionable form of trolling. In chapter 4 I argue that the combination of Michael Blome-Tillmann’s presuppositional epistemic contextualism, unjust power relations, and certain assumptions about the functions of knowledge ascriptions give rise to a host of problematic implications that are grounds for rejecting the view. In chapter 5 I argue that all current versions of conversational contextualism are unable to diagnose an important sub-class of cases of testimonial injustice.

Secondly, it helps us to develop an account of an unrecognized form of epistemic injustice that I call irrelevance injustice. Irrelevance injustice occurs either when a speaker raises an alternative that

is not taken seriously when it should be, or when a speaker raises an alternative that is taken seriously when it should not be.

Thirdly, I argue that the problems that I identify with existing forms of conversational contextualism motivate two new forms of contextualism. Chapter 7 develops a virtue epistemology inspired version of conversational contextualism that excludes the conversational moves of viciously motivated actors from influencing the standards for 'knows'. Chapter 8 develops a stakes-based view that divorces the determination of the standards for 'knows' from conversational dynamics altogether, and instead ties them to what's at morally at stake in the context of ascription. I stop short of a full endorsement of either of these views, and instead rest content with showing that they provide solutions to the problems I identify for existing forms of contextualism in this thesis whilst retaining some of their central selling points, and that as such both are worthy of further development.

Declaration

This is to certify that:

- I. This thesis comprises only my original work towards the PhD, except where stated in the preface;
- II. Due acknowledgement has been made in the text to all other material used; and
- III. This thesis is fewer than 100, 000 words in length, exclusive of table, maps, bibliographies, and appendices.

William Tuckwell

30th July 2021

Preface

Chapter 3 of this thesis draws heavily on a previously published co-authored paper with Kai Tanter.

Our contributions were equal.

Chapter 6 of this thesis draws heavily on a forthcoming single authored paper.

The publication details for each paper are as follows:

1. Tuckwell, W. & Tanter, K. (2020). 'Scorekeeping Trolls', *Thought: A Journal of Philosophy*, 9(3): 215 - 224. DOI:10.1002/tht3.464

Accepted: 20/10/2019

2. Tuckwell, W. (forthcoming). 'The Politics of Relevant Alternatives', *Hypatia: A Journal of Feminist Philosophy*

Accepted: 12/07/2021

Acknowledgements

Completing this project took a long time and involved the help of many people. I'd like to take this chance to thank some of them.

For written feedback on chapter drafts and the papers connected to them, I thank Paul Arnaud, Renee Jorgensen Bolinger, Stephanie Collins, Alex Davies, Dan Halliday, Ole Koksvik, Paul Podosky, Richard Rowland, Shawn Standefer, Weng Hong Tang, and Kai Tanter.

For proofreading an entire draft in the final week before submission, I thank Kieran McInerney.

Thanks to all who provided valuable feedback on different chapters at various conferences and seminars: the Australasian Postgraduate Philosophy Conference at the University of Queensland (2017); two instalments of the Victorian Postgraduate Philosophy Workshop (2017, 2018); the University of Tartu Work in Progress Seminar (2018); the Australasian Association of Philosophy Conference at Victoria University of Wellington (2018); the Kioloa Graduate Philosophy Conference at ANU (2018); the Epistemology Reading Group at ANU (2018); the Australian Catholic University Work in Progress Seminar (2019); the Australasian Postgraduate Philosophy Conference Online (2020); and countless talks at the weekly Tuesday evening Postgraduate Philosophy Colloquia at the University of Melbourne - these were the most fun part of the week over the last 4 years.

Most of all, I thank my two supervisors Holly Lawford-Smith and Karen Jones. They are two very different philosophers with very different styles of supervision, but they share at least one similarity: they have both been immensely supportive of me. I am hugely grateful to them both.

Finally, thanks to my wonderful girlfriend, Gina Lin. She cares very little about the semantics of knowledge ascriptions, but always shows a great deal of care for me. That has made a world of difference in helping me to see this project through to its end.

Contents

Chapter 1: Introduction	9
Chapter 2: The Functions of Knowledge Ascriptions	25
Chapter 3: Scorekeeping Trolls	42
Chapter 4: Power and Presuppositional Epistemic Contextualism	59
Chapter 5: Epistemic Contextualism or Epistemic Injustice?	85
Chapter 6: The Politics of Relevant Alternatives.....	114
Chapter 6.5: Taking Stock of Conversational Contextualism’s Problems.....	147
Chapter 7: Virtue Responsibilist Contextualism	153
Chapter 8: Making Relevant Alternatives Contextualism Moral	181
Chapter 9: Conclusion	204
Bibliography.....	206

Chapter 1: Introduction

1. Introduction

'I know that p!' declared Donald. 'But can you rule out the possibility that not-p?' replied Kamala. The standards for 'knowledge' just went up, so says the epistemic contextualist. But Donald and his buddies do not take Kamala seriously. He declares once again 'I know that p!', and he gets away with it. Contextualists often write as though when a speaker makes the type of conversational move that Kamala attempts, they put in place more demanding standards for 'knows', and that everyone has an equal ability to do this. But as the familiar situation just described illustrates, this ability is not held by all. The power that a speaker has, perhaps in virtue of occupying a particular social position, affects whether their attempts to raise the standards are seen to be successful. This furthers the interests of some - making 'knowledge', and the statuses and permissions that come with it, easier to come by - whilst working against the interests of others - making the attainment of 'knowledge' a demanding task. This is the central thought motivating this thesis.

Epistemic contextualism is the thesis that the truth-conditions of knowledge ascriptions - sentences of the form 'S knows that p' - and knowledge denials - sentences of the form 'S does not know that p' - vary with the context in which they are uttered. Accordingly, holding fixed the strength of S's epistemic position and the time of utterance, 'S knows that p' can be true when uttered in one context, yet false when uttered in another, because the strength of epistemic position required for S to satisfy 'knows' varies across the two. Speakers are said to be able to initiate a shift from a less demanding to a more demanding context by executing certain conversational moves. If S asserts 'I know that p', and I say 'but, what about this not-p alternative?', S's evidence must eliminate the alternative if they're to be able to continue truly claiming to 'know'. But contextualists have

neglected the fact that the emphasised conversational contexts are themselves embedded within broader social contexts, and that facts about these broader contexts leak into conversational contexts to affect which conversational participants can pull off the conversational moves that shift the standards for 'knows'. One feature of social context that affects speakers' conversational abilities is power. The idea that conversational abilities are influenced by power is a familiar one, but the relationship between power and the context shifting mechanisms that interest contextualists has gone largely unexplored. This unexplored relationship is one that is worthy of investigation. Knowledge ascriptions matter, and so speakers' differential abilities to affect the standards for 'knows' matter too. When Donald gets away with declaring himself a 'knower', and Kamala is unable to stop him, Donald is judged to acquire a bunch of statuses and permissions that he is not entitled to. Some (e.g., Kelp 2011) say that as a 'knower', Donald marks the point where inquiry stops. Others (e.g., Hawthorne & Stanley 2008) say he is in a strong enough position to act on what he 'knows'. Not only does this work in his favour - if inquiry has reached its end point, then the case is closed and no one has reason to keep digging to stumble across what he does not want them to find out - but it works against the interests of others - afraid that continued digging will be met with criticism, you give up and do not discover what it is in your interests to find out. The project of this thesis is to investigate some of the unexplored connections between the standards shifting mechanisms discussed by contextualists and speakers' differential conversational abilities that are determined by differences in power. The approach will be to remove the idealisation that speakers have an equal ability to shift standards and replace it with the non-ideal assumption that speakers have an unequal ability to shift standards, where this unequal ability varies with unfairly distributed power. To summarise the project in one snappy phrase, it's an attempt to theorize epistemic contextualism non-ideally.

Going non-ideal in this way has three payoffs. The first is that it allows us to see unrecognised problems with various existing forms of contextualism (chapters 3, 4, & 5). As we work through these problems with specific versions of contextualism we start to see a pattern emerge: that they are all instances of a more general problem of tying the content of ‘knows’ to conversational dynamics without placing adequate constraints on how conversational dynamics affect the content of ‘knows’. The second payoff is that it enables us to identify an unrecognised form of epistemic injustice (chapter 6). The third payoff is that the identified problems with existing versions of contextualism motivate two new forms of contextualism. By the time we get to the third payoff we will have learnt that we can’t tie the content of ‘knows’ to conversational dynamics without placing adequate constraints on how they affect the content of ‘knows’. What then should the content of ‘knows’ be tied to? All three problems identified in chapters 3, 4, & 5 involve nefarious actors making conversational moves that undermine the truth of intuitively plausible knowledge claims. The view explored in Chapter 7 takes this as motivation and develops a virtue epistemology inspired contextualism that excludes the conversational moves of nefarious actors from contributing to the content of ‘knows’. All three problems also involve various kinds of moral stakes. The view explored in Chapter 8 takes this as motivation and develops a stakes-based view that divorces the content of ‘knows’ from the conversational context altogether, and instead ties it to what’s at stake in the context. This builds on existing stakes-based versions of contextualism that take practical stakes to affect the content of ‘knows’ and suggests that various kinds of moral stakes can also affect the content of ‘knows’. I stop short of a full endorsement of either of these new forms of contextualism, and instead rest content with showing that they provide solutions to the problems facing existing forms of contextualism identified in this thesis whilst retaining some of their central selling points, and that as such both are worthy of further development.

But before we get to all of that, we need to know more about what contextualism is and why the lack of investigation into the relationship between contextualism and power is an oversight that needs correcting. That's what I'll do in this introductory chapter. The plan is as follows: in section 2 I will explain what epistemic contextualism is. There are many different versions, not all of which have the features that I've described as being part of contextualism in these first few pages. I'll identify some key theoretical choice points when it comes to picking between the different versions, and zero in on the specific version of contextualism that will be the central focus of criticism in this thesis. In section 3 I'll make an important clarificatory point about contextualism and the use-mention distinction. In section 4 I will say more about the overlooked relationship between contextualism and power. In section 5 I'll point to two (very!) recent exceptions to this oversight. Finally, I will say more about the payoffs of this investigation and provide a more detailed preview of what's to come in the rest of the thesis.

2. Contextualism

Broadly characterised, contextualism is the view that the truth-conditions of knowledge ascribing sentences – sentences of the form 'S knows that p' – and knowledge denying sentences – sentences of the form 'S does not know that p' – vary with the context in which they are uttered. Contextualism is often motivated by soliciting intuitions about pairs of cases like the following:

LOW: Hannah and Sarah are discussing whether the bank is open on Saturday. Hannah recalls, 'I've been in there on a Saturday before'. She says, 'I know that the bank is open on Saturday'. Sarah agrees.

HIGH: Hannah and Sarah are discussing whether the bank is open on Saturday. Hannah says, ‘I’ve been in there on a Saturday before’. But Sarah replies ‘Banks can change their opening hours. You don’t know that the bank is open on Saturday’. Hannah agrees (variation on McKenna 2017, variation on DeRose 1992).

Assume that Hannah is in the same epistemic position in both cases, and that the bank is open on Saturday. Intuitively, contextualists claim, both Hannah’s knowledge ascription in **LOW** and Sarah’s knowledge denial in **HIGH** are true. The contextualist claims that this is explained by the fact that the truth-conditions of knowledge claims are context-sensitive. In order for ‘S knows that p’ to be true, the epistemic standards that S must meet, in addition to having a true belief, vary with the context of ascription. Hannah’s knowledge ascription is true given the relatively lax epistemic standards that are in place in **LOW**. Sarah’s knowledge denial is also true given the more demanding epistemic standards that are in place in **HIGH**. In short, what explains the differing truth-values of the knowledge claims across **LOW** and **HIGH** is that the following thesis is true:

Contextualism: the truth-conditions of knowledge claims vary with the context of utterance.¹

That will do as a broad characterisation of contextualism, but there are various ways in which the thesis might be unpacked:

- 1) In saying that the truth-conditions of knowledge claims vary, what exactly is varying?
- 2) Which features of context determine variations in the truth-conditions of knowledge claims?

¹ Note that ‘knowledge claims’ here refers collectively to both knowledge ascriptions and knowledge denials.

I'll work through these two questions in turn to illustrate some of the different ways in which the broad contextualist thesis might be unpacked and to zero in on the specific version of contextualism that will be the target of criticism in this thesis.

There are two common ways of answering the first question. One answer that our discussion has been framed in terms of up to now says that the truth-conditions of 'S knows that p' (and 'S does not know that p') vary because the *epistemic standards* that S must meet in order for the knowledge ascription to be true (or false) vary. This view is endorsed by Keith DeRose (1995, 2009). Applied to LOW and HIGH, the explanation of our judgements about the cases is that the truth-conditions of knowledge ascriptions are context sensitive because the epistemic standard that S must meet, in addition to having a true belief, varies from context to context. Hannah's knowledge ascription is true given the epistemic standards that are operative in LOW. Sarah's knowledge denial is also true given the more demanding epistemic standards that are operative in HIGH. Let's call varieties of contextualism that give this answer to question 1) *standards contextualism*:

Standards contextualism: 'S knows that p' is true if S's epistemic position is strong enough to meet the standards that are operative in the ascriber's context, where different epistemic standards are operative in different contexts.

A second answer to 1) takes its lead from the relevant alternatives theories of knowledge that were developed by the likes of Fred Dretske (1970) and Alvin Goldman (1976), and is endorsed David Lewis (1996), Michael Blome-Tillmann (2009, 2014), and Johnathan Jenkins Ichikawa (2017a). This view says that the truth-conditions of 'S knows that p' vary because the set of not-p *alternatives* that S must rule out, in addition to having a true belief, in order for the knowledge ascription to be true vary from context to context. Applied to LOW and HIGH, the set of alternatives to the proposition that the bank is open on Saturdays that's relevant in LOW is different - it doesn't

include the possibility that the bank has changed its opening hours – to the set of alternatives that are relevant in **HIGH** – which does include the possibility that the bank has changed its opening hours. Hannah’s evidence rules out the set of relevant alternatives in **LOW**, but not those that are relevant in **HIGH**. Thus, it’s true to say that Hannah ‘knows’ in **LOW** but not in **HIGH**. Let’s call varieties of contextualism that give this answer to question 1) *relevant alternatives contextualism*:

Relevant alternatives contextualism: ‘S knows that p’ is true if S rules out all of the alternatives to p that are relevant in the ascriber’s context, where different sets of alternatives are relevant in different contexts.

It’s possible to provide a rough translation schema between the two formulations: the higher the epistemic standards, the wider the set of relevant alternatives. This translation schema is only rough because whereas standards can be subjected to a total ordering, sets of relevant alternatives cannot. Just as we can line up a group of people from shortest to tallest, epistemic standards can be lined up from least to most demanding. Much of the time we’ll be able to do this with sets of relevant alternatives too. There is a clear sense in which eliminating the alternative that the bank has changed its opening hours is less demanding than eliminating the alternative that I’m being deceived by an evil demon into thinking that the bank is open. But sometimes we won’t be able to provide an ordering of relevant alternatives. There’s no clear sense in which it’s more or less demanding to eliminate the alternative that the bank has changed its opening hours such that it’s shut all day today than it is to eliminate the alternative that the bank has changed its opening hours such that it shut early today and so won’t be open when I arrive. Rather, the demands are just different. This means that whilst it is often possible to translate the demandingness of epistemic standards into the wideness of a range of relevant alternatives, and vice versa, this is not always the case. The roughness of this translation schema doesn’t matter for what’s to come in this thesis. I will talk in terms of

standards when discussing those contextualists who cash their views out in terms of standards and in terms of relevant alternatives when discussing those contextualists who cash out their views in terms of relevant alternatives. When discussing objections that apply to both versions of contextualism I may switch between the two interchangeably. You should apply the translation schema just given at these points in the thesis. When developing my own positive views I'll opt for the relevant alternatives model.²

Moving on to question 2), which features of context determine variations in the truth-conditions of knowledge claims? That is, what determines where the contextually operative epistemic standards are set, or which alternatives are relevant? I'll describe two common answers to this question. The two answers are not incompatible with one another. Some contextualists endorse both, whilst others endorse just one or the other. According to the first answer, what determines that standards are at a particular level, or whether an alternative is relevant, is a matter of practical stakes. I'll call versions of contextualism that give this answer to question 2) *stakes contextualism*:

Stakes contextualism: the truth-conditions of knowledge claims vary with how much is practically at stake for those in the context in which the knowledge claim is made.

Different contextualists have different ideas about whose stakes matter. The most common view, defended by DeRose (1992), Cohen (1999), and McKenna (2011), says that it is the stakes of the ascriber and their interlocutors that matter. On this view, the more important it is to the ascriber and their interlocutors that a proposition is true, the more demanding the standards, or the more alternatives that need to be ruled out, in order for a claim to 'know' the proposition to be true. If it doesn't matter very much whether the bank is open, then the possibility that the bank has changed

² See Johnathan Schaffer (2005) and Ichikawa (2017a) for arguments why the relevant alternatives model is superior.

its opening hours is not relevant. If it matters a great deal whether the bank is open, then the possibility that the bank has changed its opening hours is relevant. Stakes views are not the primary target of criticism in this thesis, but a version of the stakes view makes an appearance in chapter 8 where I explore the prospects of a new theory of stakes contextualism.

A second answer to question 2) says that the dynamics of the conversation in which a knowledge claim is made determine its truth-conditions. McKenna (2017) has aptly described versions of contextualism that give this answer to question 2) *conversational contextualism*:

Conversational contextualism: the truth-conditions of knowledge claims depend upon the dynamics of the conversation in which the knowledge claim is made.

Drawing heavily on McKenna's (2017) description, the view is as follows: speakers make various conversational moves that are informed by their intentions, expectations, purposes, and presuppositions. On Friday Hannah says that she 'knows' that the bank is open on Saturday as she has been there on a previous Saturday. Sarah replies that they need to get to the bank before the weekend is out and asks if Hannah has seen the latest opening hours. She points out that 'banks do change their opening hours'. Hannah concedes that she cannot rule this out. What we see then is that as the conversation develops, the truth-conditions for the knowledge claim change, i.e., the operative epistemic standard, or the set of relevant alternatives change. At the beginning of the conversation the alternative in which the bank has changed its opening hours is not relevant. It hasn't occurred to anyone, no one has mentioned it, etc. By the end of the conversation the alternative is relevant. Everyone accepts that it is relevant, that it needs to be considered, etc. The various conversational moves determine the truth-conditions of knowledge ascriptions. This is the view (or more precisely, particular versions of it) that will be the target in the critical part of this thesis in chapters 3, 4, and 5.

We can also describe the process by which the truth-conditions of knowledge claims change in terms of a changing conversational score (Lewis 1979, Langton 2018). It is worth also describing conversational contextualism in these terms as it will come up at a few points later in the thesis. Conversational score is an abstract register of the moves that are made in a conversation which has normative force in that it determines what further moves speakers are permitted to make at a given time. One function of score is to help determine the contents of context-sensitive expressions. The conversational score, having registered the different conversational moves that speakers have made up until a given time, will determine what relation is picked out by ‘knows’. As further conversational moves are made, the score will update, and the relation that is picked out by ‘knows’ changes too.³

Prominent contextualists who hold some version of conversational contextualism include Lewis (1996), Stewart Cohen (1988), DeRose (1995, 2009)⁴, Blome-Tillmann (2009, 2014), and Ichikawa (2017a).⁵ It’s a theoretical option to combine either of standards contextualism or relative alternatives contextualism with either or both of stakes contextualism or conversational contextualism.

³ Here I follow Langton’s (2018) understanding of conversational score. She notes that Lewis provides two ways of understanding conversational score. One is the way that I’ve just described. Another way is in terms of the psychological states of conversational participants. Participants alter their beliefs and attitudes as various conversational moves are made. They then continue the conversation with these shared assumptions in place. Langton reserves the term ‘common ground’ for this latter understanding.

⁴ Notice that both DeRose and Cohen think that both stakes and conversational dynamics influence variations of the truth-conditions of knowledge claims.

⁵As noted above, DeRose (1992) thinks that stakes play an important role in determining the contextually operative epistemic standards. Lewis also takes stakes to play a role in determining whether an alternative is relevant by determining what degree of belief a subject has in a possibility makes it relevant. When error would be especially disastrous a low degree of belief makes an alternative relevant, and few possibilities are properly ignored (1996: 556). Blome-Tillmann (2009: 262) thinks that which alternatives are relevant is a matter of which alternatives are ‘taken seriously’, and that stakes influence which possibilities speakers take seriously.

3. Use and Mention

When discussing epistemic contextualism it is important to keep in mind the distinction between use and mention. To use a word is to employ it when speaking or writing in order to talk about the thing that the word refers to. For example, in uttering ‘Melbourne sees four seasons in one day’ I use the word ‘Melbourne’ to talk about the city’s famously changeable weather. To mention a word is to talk about the word itself, rather than the thing that the word refers to. For example, in uttering ‘‘Melbourne’ starts with M’ I mention the word ‘Melbourne’ to talk about a feature of the name of the city rather than the city itself.

Statements of epistemic contextualism mention rather than use the word ‘knows’. When contextualists say that Hannah ‘knows’ relative to the low standards operative in **LOW** and that Hannah does not ‘know’ relative to the higher standards in **HIGH**, sometimes satisfying the predicate ‘knows’ requires meeting a weaker epistemic condition and sometimes it requires satisfying a stronger one. But this is not to say that whether one knows in the sense of possessing the metaphysical property of knowledge depends on what conversational moves are made.

Failure to acknowledge this distinction can lead to confusion, and so it’s important to keep it clear. For example, some people object that making certain conversational moves shouldn’t make it the case that someone loses knowledge (e.g., Goldman 2007: 17). The contextualist response is that it doesn’t. The subject still has the property that they had before the conversational move was made, it’s just that the word ‘knows’ expresses a stronger relation in **HIGH** than in does in **LOW**. Only by keeping the distinction between use and mention clear can we avoid such confusion.

A standard convention is to put quotation marks around a word to designate that it is being mentioned. When talking about the thing that epistemic contextualism is a thesis about, I will either talk about knowledge ascriptions (to be read as ‘a sentence expressing that someone satisfies

‘knows’’) and knowledge denials (to be read as ‘a sentence expressing that someone fails to satisfy ‘knows’’) or follow the convention of putting quotation marks around ‘knows’ as much as is possible.⁶

4. Power

Conversational contextualism emphasises the role of conversational dynamics in the determination of the truth-conditions of ‘knowledge’ claims. But conversational contextualists have failed to recognise the fact that conversational dynamics and speakers’ abilities are affected by power. The relationship between power and conversational abilities has been much discussed in debates about silencing (see, e.g., Langton 1993, Langton & Hornsby 1998, Langton & West 1999, Maitra 2009, McGowan 2009, Davies 2016, Hesni 2018), subordinating speech (see, e.g., Langton 1993, Maitra 2012, West 2012), and epistemic injustice (see, e.g., Jones 2002, Fricker 2007, Dotson 2011, Pohlhaus 2012, Medina 2013), but the relationship between power and the context shifting mechanisms that interest contextualists has gone largely unexplored.

The idea that speakers have equal abilities to affect the content of ‘knows’ is common amongst conversational contextualists. Lewis famously (1996: 559) claims that if attention is paid to some not-p alternative, then it is thereby relevant. DeRose (2004: 3, 2009: 133) claims that every speaker’s personally indicated standards register in determining the truth-conditions of ‘knows’, and so seems to assume that there is an even linguistic playing field. We will see in chapter 4 that the success of Blome-Tillmann’s (2009, 2014) contextualism depends upon fair co-operation among conversational participants. But none of these assumptions hold in non-idealised conversations.

⁶ See Ichikawa (2017a: 1 – 3) and (2017b: 3 – 4) for more in-depth discussion on the significance of the use-mention distinction to debates of epistemic contextualism.

Sometimes a speaker who is lacking in power draws attention to an alternative but is ignored by those who are more powerful. There is not an even linguistic playing field. Rather, there are asymmetries in speakers' linguistic capabilities that are sometimes a function of asymmetries in power. Speakers are not always co-operative. And, even when they are, there is no guarantee that the terms of co-operation are fair. Chapters 3, 4, and 5 show how unrecognised problems for existing versions of contextualism arise when the falsity of these assumptions is interrogated.

These observations are important. Firstly, insofar as contextualism is supposed to be a descriptive thesis, then descriptive adequacy demands that it pay attention to non-ideal aspects of conversational context. Secondly, knowledge ascriptions matter. As we'll see in detail in the next chapter, being ascribed 'knowledge' gives one a range of statuses and permissions. If the standards for 'knows' can be manipulated by conversational moves, and what conversational moves one can make are influenced by how much power one has, then who gets to count as 'knowing' and how much epistemic labour one needs to carry out prior to being ascribed 'knowledge' will also be a function of these power asymmetries. This thesis removes the idealisations about equal conversational abilities made by conversational contextualists, replaces them with more realistic non-ideal assumptions about unequal conversational abilities, and investigates what the implications are for conversational contextualism.

5. Some recent exceptions

When I began writing this thesis in 2017 there was, to my knowledge, only one notable discussion of the relationship between contextualism and power. That came from Evelyn Brister (2009, 2017). My plan was to pick up from where she left off. During the time I've been writing, one more philosopher - Johnathan Jenkins Ichikawa (2020) - has also picked up on the importance and

fruitfulness of the relationship between various aspects of contextualism and power. I'll briefly explain the relationship between my project and theirs here. I'll provide a more sustained engagement with their work at various points throughout the thesis.

Briester (2009, 2017) has expressed concern that contextualism is too skeptic friendly. As I've said, some contextualists take it that when a speaker raises a skeptical hypothesis to salience, higher standards for 'knows' are put in place. Briester is concerned that this unfairly concentrates epistemic power in the hands of skeptics in a way that 'serves epistemically conservative interests, and severs the connection between knowledge and action' (2009: 683). Because skeptics can always invoke high burdens of proof for new and novel claims, they always have the option to shut down inquiry, and to bully and silence inquirers as they see fit. Skeptic friendly versions of contextualism thus have a tendency towards conservatism that favours those whose interests are well served by the status quo, and works against those whose interests are not well served. Here is a simple example: suppose that a scientist claims that second-hand tobacco smoke has serious negative health impacts. Tobacco companies respond by introducing some hard to eliminate error possibility. Claims to 'know' that second-hand tobacco smoke cause serious negative health impacts are thus undermined. This is epistemically conservative in that knowledge production is blocked. This serves the status quo by enabling those who sell tobacco to continue doing so without having to battle the reputational damage that would be done by 'knowledge' of its harmful effects. This works against the interests of those who smoke and those who inhale second-hand smoke. What's more, Briester claims, because action is often based on what we think we know, skeptical challenges can undercut our ability to act on the basis of justified beliefs (2009: 683). It's more difficult to justify placing warning labels of negative health impacts on cigarette packets if claims to 'know' about the adverse health impacts of smoking have been undermined by the raising of skeptical error possibilities. In a passage that nicely captures Briester's concerns, she writes:

[S]kepticism is a kind of power grab that puts other epistemic agents on the defensive and at a disadvantage. While some might believe that knowledge is power, the skeptic recognizes that *denying* knowledge is power. (2009: 682).

This a great insight. I'll explore a variety of ways in which making skeptical conversational moves can be advantageous for the speaker, whilst proving disadvantageous for their interlocutors, across chapters 3, 4, 5, and 6.

Ichikawa (2020) has recently made use of the insight that there are unjust asymmetries in people's abilities to influence and determine conversationally operative standards for 'knows' in order to motivate the category of *contextual injustice*. Contextual injustice occurs when the standards for 'knows' are prejudicially raised to an inappropriate level, and a testifier is wrongfully denied 'knowledge' as a result (2020: 14). Ichikawa thinks that the knowledge denials that are involved in contextual injustice are literally true, but are nevertheless unjust because the standards the subject of the knowledge denial is being evaluated against are inappropriate to the context at hand. Contextual injustice bears some similarities to the notion of *irrelevance injustice* that I identify in chapter 6, but as we shall see the two are importantly different.

6. A preview of the payoffs

In this chapter I have described epistemic contextualism, drawn attention to an oversight in the existing contextualist literature, and explained why it is important to correct for this oversight. It has to do with knowledge ascriptions having important functions. Chapter 2 is devoted to explaining the importance of knowledge ascriptions at greater length, and thus provides further motivation for

a thesis long investigation into a particular theory of knowledge ascriptions. What do we stand to get out of investigating the connection between conversational contextualism and power?

I'll answer this question by way of previewing what's to come in the rest of the thesis. In chapter 3 and 4 we see new problems with two specific versions of conversational contextualism, Keith DeRose's (2004, 2009) gap view (chapter 3) and Michael Blome-Tillmann's (2009, 2014) presuppositional epistemic contextualism (chapter 4). In chapters 5 and 6 we learn of the connections between epistemic contextualism and epistemic injustice. In chapter 5 I'll argue that current versions of conversational contextualism can't vindicate our intuitions that certain paradigm cases of testimonial injustice are in fact cases of testimonial injustice. In chapter 6 I'll argue that conversational contextualism has the resources to illuminate an unrecognised form of epistemic injustice that I call irrelevance injustice. In chapters 7 and 8 we'll see two new versions of epistemic contextualism that are well placed to deal with the problems adduced for conversational contextualism in chapters 3, 4, and 5: virtue responsibilist contextualism (chapter 7) and moral stakes contextualism (chapter 8). I'll argue that both of these views can handle the problems that I identify with extant versions of conversational contextualism, whilst retaining some of their central selling points. I'll also respond to some objections to each of the views. I'll stop short of endorsing either of the views outright - much more work needs to be done before I can make that commitment! Rather, I rest content with showing that both views hold enough promise to warrant further development.

Chapter 2: The Functions of Knowledge Ascriptions

1. Introduction

Suppose that Brett says, ‘Dan knows that the pub is shut’. Yes, Brett has just told us something about Dan’s mental state. But he has done more than this. He has also marked Dan as a reliable informant; as someone that we can go to for information about whether the pub is open (Craig 1990). Brett has also signalled that inquiry stops with Dan. The pub is shut; no more questions need be asked (Kelp 2011). And Brett tags Dan as someone we cannot criticize for asserting (Williamson 2000), acting on (Hawthorne & Stanley 2008), or believing (Williamson 2000) that the pub is shut, at least not on epistemic grounds. Being in the ‘know’, Dan is permitted to do these things.⁷

All of that is to say, knowledge ascriptions – sentences of the form ‘S knows that p’ – have been said to play a variety of functions. Whichever of these views is correct, it is important both that our semantic theories of knowledge ascriptions and our practices of ascribing and denying knowledge are accurate. Getting things wrong is costly. If Craig is right, and if Brett has made a mistake, then Dan has been marked as reliable when he is unreliable, and he will receive trust that he does not deserve. Making the opposite mistake of marking the reliable as unreliable will mean that the reliable miss out on trust that they are entitled to. If Kelp is right, making a mistake about whether someone is a ‘knower’ is to make a mistake about whether inquiry has reached an adequate end point, and this is costly too. Declare prematurely, and you risk missing out on a drink at your

⁷ For readers outside of an Australian context some extra details might be required to understand the examples I’m using here. ‘Dan’ refers to Daniel Andrews who is the Premier of the Australian state of Victoria. ‘Brett’ refers to Brett Sutton who is Victoria’s Chief Health Officer. Jointly, they are presiding over the Victorian response to the Covid-19 pandemic. Andrews’ pub comments both devastated and inspired Victorians as it both meant that they missed out on going to the pub, but also led to the creation of songs, t-shirts, and gigantic murals devoted to Dan and his pub talk.

favorite pub. Take too long, and you'll waste valuable time that could be used to find some non-drinking related activity to keep you occupied whilst you wait out the lockdown. If Williamson, and Hawthorne & Stanley are right, mistakes lead us to judge the prohibited to be permitted. If, for example, Brett has made a mistake about Dan, and that mistake goes unrecognized, Dan's asserting that 'the pub is shut' will go uncriticized.

This chapter surveys these various proposals about the function(s) of knowledge ascriptions in order to help us to get clear about their importance. This will put us in a better position to see what's at stake when our semantic theories of knowledge ascriptions and our knowledge ascribing and denying practices go wrong, as I'll argue that contextualism does when it's theorized against the backdrop of real world power relations. Combining contextualism, real world power relations, and different assumptions about the function of knowledge ascriptions gives rise to an abundance of problematic implications - epistemic, moral, and political - that are, I'll argue over the course of chapters 3, 4, and 5, reasons to reject extant versions of conversational contextualism.

I do not take a stand on whether any of these proposals correctly identifies a function of knowledge ascriptions: they all enjoy some intuitive plausibility. Rather, I aim to show that knowledge ascriptions are important whichever of the views turns out to be correct. This means that whatever the function of knowledge ascriptions turns out to be, the fact that conversational contextualist theories of knowledge ascriptions go wrong is of concern.

Section 2 describes Edward Craig's (1990) proposal that knowledge ascriptions certify reliable informants. Section 3 describes Christoph Kelp's (2011) proposal that knowledge ascriptions signal the end of inquiry. Section 4 describes the connection between knowledge ascriptions and the norms of assertion (Williamson 2000), action (Hawthorne & Stanley 2008), and belief (Williamson 2000). Section 5 describes contextualist treatments of the three knowledge norms described in

section 4. Section 6 considers the connections between knowledge ascriptions and epistemic injustice, providing further evidence of the importance of knowledge ascriptions. Section 7 summarises the chapter and anticipates what is to come in those that follow.

2. Knowledge Ascriptions Certify Reliable Informants

Discussions of the function of knowledge ascriptions commonly start with Edward Craig's 1990 book *Knowledge and the State of Nature*. Craig provides us with both a methodology for uncovering the function of knowledge ascriptions and a hypothesis about the function of knowledge ascriptions. He describes his methodology as follows:

We take some prima facie plausible hypothesis about what the concept of knowledge does for us, what its role in our life might be, and then ask what a concept having that role would be like, what conditions would govern its application... (1990: 2-3).

Putting this methodology to work, Craig develops a genealogical story that begins in an epistemic state of nature and arrives at the hypothesis that the function of knowledge ascriptions is it to certify sufficiently reliable informants.

Craig's genealogical story starts out in a community of our imaginary ancestors. He notes that they have certain needs and interests. Firstly, they have an interest in gaining information about the world. Sometimes they're able to obtain information by themselves through perception, introspection, and the like. But because of their various limitations, such as their cognitive limitations or because they only have a finite amount of time, there will be times when they cannot obtain information by themselves. Instead, they will need to appeal to others for information. For example, they might ask others if certain foods are safe to eat. Secondly, they have an interest in

being able to distinguish reliable from unreliable sources of information. Craig then goes on to suggest that for creatures with this pair of interests it would be useful to have some way of identifying and sorting those who are reliable from those who are unreliable sources of information. Call this way of identifying reliable sources of information the practice of ascribing knowledge*, an ancestor of our current practice of ascribing knowledge. How reliable does a source of information need to be in order to count as a good source? Craig thinks that the answer to this question would have initially had a subjective answer: the source would need to be as reliable as first-person needs require. While the notion of a good source of information might have started out life with this subjective character, over time it would naturally evolve to take into account the concerns of others as well. This is because humans are information gathering and information sharing creatures who want others to consider their information gathering needs. In response to these needs, over time the practice of ascribing knowledge* becomes objectified to become our current practice of ascribing knowledge. A knowledge ascription thus comes to be used to certify agents as reliable enough for a wide variety of people with diverse goals and interests. Given that people with this diversity of goals and interests will require differing degrees of reliability, the practice of setting the standards for sufficient reliability fairly high becomes established. We thus arrive at a practice in which knowledge ascriptions certify people as reliable enough for a diverse range of people with different interests and goals to draw upon.

Craig's proposal is a popular one. Daniel Greco (2007, 2008), Miranda Fricker (2008), Martin Kusch (2009), and Duncan Pritchard (2009) have all appealed to Craig's hypothesis to support various epistemological views. But it also has numerous detractors, some of which have an alternative story about the function of knowledge ascriptions. It is to these detractors that we turn to next.

3. Knowledge Ascriptions Signal the End of Inquiry

Christoph Kelp (2011) has argued that Craig misidentifies the function of knowledge ascriptions. Using Craig's methodology, Kelp arrives at a different hypothesis: that the function of knowledge ascriptions is to flag when agents can adequately terminate inquiry into a given question (2011: 62).

Why think this? Kelp says that not only do we have a need to be able to identify reliable informants, but we also have a need to be able to evaluate agents' inquiries. Kelp (2011: 62) argues that it is this need that the function of knowledge ascriptions is responsive to. Inquiry into particular questions can be open ended. What does it take to know that the pub is shut? Must I be told by Dan? Must I check the news to ensure I'm up to date with the latest coronavirus restrictions? Must I go down to the pub and bang on the door until an angry pub landlord gives me an answer? But given our concerns (e.g., our practical and theoretical interests) and capacities (e.g., limited time), we need a way to certify that inquiry has been extensive enough. Knowledge ascriptions give us a way to do this: when we say 'S knows that p' we mark a point where inquiry can be stopped. Similar proposals have also been made by Klemmens Kappel (2010) and Patrick Rysiew (2012).

Michael Hannon (2015: 858 - 859) thinks that while Kelp (and Kappel and Rysiew) is surely right that knowledge ascriptions function as inquiry stoppers, this is not incompatible with them having the function of identifying reliable informants. Rather, Hannon suggests these are two sides of the same coin. A reasonable way to bring inquiry to an end is to identify someone as a sufficiently reliable informant on the question at hand. If Dan is reliable as to whether the pub is shut, then I can take his word for it. The question is settled. Inquiring further would be both a waste of time and cognitive resources.

4. Knowledge Ascriptions Track Epistemic Norms

Some have argued that knowledge is the norm of assertion (Williamson 2000), action (Hawthorne & Stanley 2008), and belief (Williamson 2000). The norms are knowledge norms in the sense that they say that knowledge is necessary and/or sufficient for epistemically permissible assertion, action, and belief. These norms are influential, but are not without detractors. The point here is not to defend them, but rather to draw attention to some of the potentially important functions of knowledge ascriptions.

It's also important to note, as Michael Hannon (2015: 863) has, that the philosophers who argue for these norms don't explicitly make claims about the functional role of knowledge ascriptions. But because their views are supported by appeals to knowledge ascriptions and denials, they are assuming that knowledge-talk is tracking these norms. For assertion and action, this is because these norms line up with assessments of assertions and actions in ordinary language. For belief, this is because there is a parallel between belief and assertion, and because knowledge-talk lines up with assessment of assertions in ordinary language.

I'll state these norms in their biconditional form (Hawthorne & Stanley (2008), for example, endorse a biconditional version of the knowledge norm of action). This is both for simplicity of presentation and because it's these versions of the norms that will be discussed later in this thesis.⁸

I'll work through each of these norms in turn, stating both the norm and some motivations.

In all three of the norms 'permitted' is to be given an epistemic reading. This means that nothing relating to S's epistemic position can make it impermissible for S to, e.g., assert that p. But, this is

⁸ One might adopt 'one-direction' versions of the norms according to which knowledge is either necessary or sufficient. Williamson (2000), for example, endorses the necessity direction of the knowledge norm of assertion. Fantl & McGrath (2009) endorse the sufficiency direction of the knowledge norm of action. In the course of discussing the motivation for the bi-conditional versions of the norms we will also see the motivations for the one-direction versions of the norms.

compatible with other things making it impermissible. For example, S asserting that p might cause the listener to be upset.

4.1. Assertion

ASSERTION: S is permitted to assert that p iff S knows that p.

ASSERTION gets its motivation from its ability to explain various conversational patterns and bits of linguistic data. The necessity direction is supported by its ability to explain our practices of challenging people's assertions. It's appropriate to challenge an asserter by asking 'how do you know that?' The necessity direction of ASSERTION also explains why it's not acceptable for the asserter to respond by claiming that they did not claim to know what they asserted (Williamson 2000: 252) (Unger 1975: 250 - 65). The sufficiency direction of ASSERTION explains why it is acceptable for the asserter to respond by claiming that they know what they asserted (DeRose (2009: 93) endorses the biconditional version of ASSERTION).

Data from lottery assertions provides further support for the necessity direction of ASSERTION. Supposing that a lottery is sufficiently large, there is a very high probability that any given ticket will lose. Yet, many find it inappropriate to flat out assert that a given lottery ticket will lose. ASSERTION provides a ready explanation: it is inappropriate to assert that a given ticket will lose, because one does not know that the ticket will lose (Williamson 2000: 246 - 252) (Hawthorne 2004: 21 - 23).

4.2. Action

ACTION: S is permitted to use p as a premise in her practical reasoning iff S knows that p.

ACTION gets its motivation from its ability to explain our practices of criticizing and defending actions. We take it to be appropriate to criticize someone's actions if they act in situations where they do not know. 'You shouldn't have gone down this street, since you didn't know the restaurant was there' (Hawthorne & Stanley 2008: 571). We also think that an agent may appropriately defend herself by appealing to what she knows. 'I know that the restaurant is on this street, so I didn't do anything wrong in driving down it.'

ACTION's ability to explain data from lottery assertions can also be marshalled in its support. Consider the following scenario taken from Hawthorne & Stanley (2008: 572). Suppose that you are offered a cent for a lottery ticket that cost you a dollar. There are 10, 000 tickets in the draw and the prize ticket wins 5, 000 dollars. Suppose that you reason as follows:

I will lose the lottery

If I keep the ticket, I will get nothing

If I sell the ticket, I will get a cent

So, I ought to sell the ticket.

Hawthorne & Stanley (2008: 572) think that many will find this pattern of practical reasoning absurd. **ACTION** provides a natural explanation of its absurdity: one cannot act on the belief that one will lose in the lottery since one does not know that one will lose.

4.3. Belief

BELIEF: S is permitted to believe that p iff S knows that p.

A common argument for **BELIEF** starts out by identifying a parallel between belief and assertion according to which belief amounts to an inner assertion. For example, Timothy Williamson writes:

It is plausible, nevertheless, that occurrently believing that p stands to asserting that p as the inner stands to the outer. If so, the knowledge rule for assertion corresponds to the norm that one should believe p only if one knows p. Given that norm, it is not reasonable to believe p when one knows one does not know that p (2000: 255 - 256).

Those who find this parallel plausible and who also accept **ASSERTION** have reason to accept **BELIEF**. Knowledge ascriptions track the norm of belief in that when we say ‘S knows that p’ we not only mark S as someone who is permitted to assert that p, but also as someone who is permitted to believe that p.

4.4. Objections to the Knowledge Norms

Let’s turn to consider some of the objections that have been levelled against the knowledge norms. One way of objecting to knowledge being necessary for assertion, action, and belief, has been to argue that there are cases in which asserting that p, relying on p in practical reasoning, and believing that p are all permissible even though one doesn’t know that p. For example, in cases of luckily justified true beliefs such as Gettier cases, asserting, acting on, and believing that p is all permissible (see, e.g., Brown 2008a). Williamson (2000: 256) and Hawthorne and Stanley (2008: 586) have responded by saying that these purported counterexamples are really just excusable violations of the norms.

Objections have also been levelled against the sufficiency direction of ACTION. Jessica Brown (2008b) provides the example of a surgeon who double-checks her notes before operating on the diseased left kidney of a patient even though she knows that it is the left kidney which is the diseased one. A nurse is explaining what's going on to a group of medical students who are observing the procedure. She says, 'of course the Dr knows that it's the left kidney that's diseased, but she shouldn't set to work before double-checking the patients notes.' Jonathan Jenkins Ichikawa (2012: 49) argues that cases in which intuitively a subject knows, but in which they are not intuitively warranted in acting, don't count against ACTION. He claims that Brown's proposed counterexample is unsuccessful because while ACTION says that you can rely on p as a reason for action iff you know that p, this doesn't mean that knowledge that p is sufficient for action. Perhaps under the current circumstances a higher epistemic status than knowledge is required. Perhaps the surgeon must know that she knows that it is the left kidney that's diseased before operating, and the surgeon doesn't know this (2012: 51).

5. Contextualism and the Knowledge Norms

Regardless of whether it is the proponents or the detractors of the knowledge norms who are correct, contextualists have been greatly interested in them. Michael Blome-Tillmann (2013) claims that the norms have sufficient plausibility for their incompatibility with contextualism to count against contextualism. He thus goes to great lengths to respond to John Hawthorne's (2004: 87) arguments that contextualism is incompatible with the knowledge norms. Keith DeRose (2009: 247 - 251) is in agreement with Blome-Tillmann with respect to ACTION. He has gone further with respect to ASSERTION, claiming that it provides a powerful argument for contextualism (DeRose 2002). It is because contextualists have shown so much interest in the knowledge norms that I

devote so much attention to them throughout the thesis. Across chapters 3, 4, and 5 we'll see that the combination of contextualism, real world power relations, and the knowledge norms gives rise to an abundance of concerning implications; implications that offend not only our epistemic intuitions, but our moral and political intuitions too.

As I've described them up to now, the three knowledge norms say that whatever context *S* finds herself in, it is permissible for *S* to assert, act on, and believe that *p* iff she satisfies the content of 'knows' that is operative in the context of this thesis chapter. But, contextualism says that 'knows' picks out different contents in different contexts. This means that having 'knows' as it figures in this thesis chapter set the standards of permissibility is arbitrary from the perspective of the contextualist. The knowledge norms thus need to be reformulated to reflect the context-sensitivity of 'knows'. Cohen (2004: 486), DeRose (2009: 99), and Blome-Tillmann (2013) all suggest the following reformulations:

ASSERTION RELATIVIZED: *S* is permitted to assert that *p* iff *S* satisfies 'knows that *p*' in *S*'s context.

ACTION RELATIVIZED: *S* is permitted to use *p* as a premise in practical reasoning iff *S* satisfies 'knows that *p*' in *S*'s context.

BELIEF RELATIVIZED: *S* is permitted to believe that *p* iff *S* satisfies 'knows that *p*' in *S*'s context.

Now to DeRose's (2002) argument from **ASSERTION RELATIVIZED** for contextualism. The argument starts out with the observation that what you are permitted to assert is often judged to depend on your practical situation and what is at stake. When it matters a great deal to you whether a proposition is true, the standards for permissibly asserting the proposition are much higher than

when nothing much is at stake. For example, suppose you're taking a journey by public transport to go to visit an elderly relative. Now, coronavirus is an ever-present threat these days. Given your exposure to other people on your route and the fact that the elderly are particularly vulnerable to the virus, the standards for permissibly asserting that 'my hands are clean' are high. Compare this to the situation in which you're on your way to take the rubbish out. In this situation the standards for permissibly asserting 'my hands are clean' are more relaxed. This observation, together with **ASSERTION RELATIVIZED**, yields the conclusion that 'knows' is context sensitive. If whether one is permitted to assert varies with context, and being permitted to assert depends on whether you know (i.e., **ASSERTION RELATIVIZED** is true), then 'knows' varies with context too (DeRose 2002: 187).

6. Knowledge Ascriptions and Epistemic Injustice

Our epistemic practices are problematically influenced by prejudice and power in ways that give rise to epistemic injustice: injustices that consist in wronging agents in their capacity as a knower (Fricker 2007:1) or their capacity as an epistemic subject (Gerken 2019: 2). The practice of ascribing and denying knowledge is one place where the problematic influence of prejudice and power takes hold. Given the variety of socially valuable functions of knowledge ascriptions we should expect there to be interesting and important connections between knowledge ascriptions and epistemic injustice that make investigating theories of knowledge ascriptions worthwhile. Surprisingly little has been said about this connection. Here are three examples, each of which are investigated at length later in the thesis.

Knowledge ascriptions and the harms of epistemic injustice. Paying attention to the social functions of knowledge ascriptions can shed light on the harms wrought by epistemic injustice. For example,

such attention can elucidate unrecognized harms of testimonial injustice, the injustice that occurs when a speaker is accorded less credibility than they deserve when testifying because of the operation of an identity prejudice in the audience (Fricker 2007: 27). I'll argue in chapter 5 that one way that a speaker can be accorded less credibility than they deserve is by a hearer responding to their testimony by wrongfully denying that they 'know'.⁹ When it is denied that a subject 'knows', they are denied the socially valuable statuses and permissions that come with being ascribed knowledge. One is perceived not to have the status as a reliable informant, a status that one should have. One is thereby not accorded the esteem that comes with such a status. The speaker is marked as someone whose inquiry must continue, rather than as at a place where inquiry can stop. This means that they incur the cost of having to expend more epistemic labor than they should have to. The speaker is not taken to be permitted to assert, act on, or believe the proposition that they testify to. If they do, they open themselves up to criticism.

These various costs and harms are not mentioned in current discussions about the harms of testimonial injustice, which focus instead on whether we should think of the victim of testimonial injustice as someone who is *epistemically objectified* by being treated as someone who is completely lacking in epistemic agency (see, e.g., Fricker 2007: 132, McGlynn forthcoming), or as someone who is *epistemically othered* by being treated as someone whose epistemic agency is only partially recognized (see e.g. Pohlhaus 2014, Davis 2016). The existing discussion of the harms of testimonial injustice are far more complex than I can hope (or need) to do justice to here, but the important point is that connection between knowledge ascriptions and the harms of epistemic

⁹ On Fricker's account a speaker suffers testimonial injustice when they are wrongfully denied as possessing some epistemic property such as knowledge or some particular level of justification, not when they are wrongfully denied satisfying 'knows'. I'll argue in chapter 5 that testimonial injustice should also be understood as applying to cases where a subject is wrongfully denied satisfying 'knows'.

injustices such as testimonial injustice is missing from extant discussions. Given their socially valuable functions, this is something worth investigating further.

Theories of knowledge ascriptions and difficulties diagnosing epistemic injustice. This second connection is not so much to do with the function of knowledge ascriptions, but rather to do with how different theories of knowledge ascriptions might help or hinder us in diagnosing epistemic injustice. Mikkel Gerken (2017, 2019) has argued that pragmatic encroachment theories of knowledge and knowledge ascriptions obscure, and perhaps even sanction, epistemic injustice.¹⁰ Pragmatic encroachment says that knowledge and knowledge ascriptions depend in part on practical factors such as how much is at stake for the subject of knowledge. The more that is at stake, the more evidence required for a subject's belief to be justified, the harder it is for a subject to know. Pragmatic encroachment differs from contextualism in important ways. Firstly, whereas contextualism is solely a linguistic thesis about the meanings of knowledge claims, pragmatic encroachment is a substantive metaphysical claim about the nature of knowledge itself. It says that whether or not a subject knows depends upon their practical situation. Secondly, whereas contextualism says that the meanings of knowledge claims vary with facts about the context in which the claim is made, pragmatic encroachment says that whether or a subject knows depends upon facts about the subject's own context; in particular, their practical situation. Thirdly, whereas contextualism allows that an ascription and a denial of 'knowledge' of the same proposition, to the same subject, at the same time can both be true when uttered in different contexts, pragmatic encroachment does not allow for this. Instead, it says that the truth-values of knowledge ascriptions cannot vary across contexts of ascription holding fixed time, subject, and proposition. Pragmatic encroachment is thus a version of invariantism, a competitor to contextualism. Hence why Jason

¹⁰ Amia Srinivasan (2016) notices a similar problem.

Stanley (2005) labels his version of pragmatic encroachment ‘interest relative invariantism’.¹¹

Gerken (2019: 7) provides the following case as an illustration of the difficulty that pragmatic encroachment has with diagnosing epistemic injustice:

After years of pursuing employment and many rejected applications, Brooke has been hired as a secretary in a small company. Unfortunately, her supervisor has started to sexually harass her. Brooke considers reporting that she’s being sexually harassed. However, she does not have hard evidence, and she suspects that she might be laid off if she testifies against her supervisor. Although she desperately needs to keep the job, Brooke asserts that she has been sexually harassed by her supervisor.

Assume that because of her gender Brooke’s audience judge her to be too emotional to know whether her supervisor sexually harassed her or was just flirting; due to identity prejudice they regard her as a non-knower. It seems to Gerken, and to me too, that Brooke suffers testimonial injustice. She is judged to be in a weaker epistemic position than she is in fact in because of prejudice. But, given Brooke’s high stakes situation, pragmatic encroachment theories say that Brooke does not know and so her audience judging that she does not know is not an undervaluation of her epistemic position. This means that the pragmatic encroacher cannot diagnose the case as an instance of epistemic injustice. (Or, it at least remains to be shown how the pragmatic encroacher can diagnose epistemic injustice in these cases). We should expect pragmatic encroachment theories to have trouble diagnosing cases of testimonial injustice fairly often. Those who lack social and economic privilege frequently find themselves in high stakes situations in virtue of their lack of privilege. When people who are in this situation are treated as non-knowers, they cannot be said to

¹¹ See Kim (2017) for a helpful overview of pragmatic encroachment views.

be victims of epistemic injustice by the lights of pragmatic encroachment. To foreshadow the argument to come in chapter 5, things are even worse for epistemic contextualism. As I will argue, not only does epistemic contextualism have difficulty diagnosing cases of epistemic injustice, but it also provides the resources for hearers to perpetrate testimonial injustice like harms and get away with it.

Knowledge ascriptions and new forms of epistemic injustice. The third connection is that investigating the connection between knowledge ascriptions and epistemic injustice has the potential to reveal unrecognized forms of epistemic injustice. We will see this potential realized when I develop a theory of an unrecognized form of epistemic injustice that emerges out of an investigation into contextualist theories of knowledge ascriptions in chapter 6. To anticipate, the basic idea will be that an injustice lies in the fact that differentially socially situated agents have asymmetric power to affect the conditions under which knowledge is ascribed, and this in turn affects the ease and difficulty with which differentially socially situated agents are ascribed knowledge.

7. Summary and Preview

We've seen that knowledge ascriptions play socially valuable functions. We've also seen that there are significant connections between knowledge ascriptions and epistemic injustice. Across chapters 3, 4, 5, and 6 we'll see an abundance of epistemically, morally, and politically concerning implications that arise from combining current versions of conversational contextualism with real world power relations and different assumptions about the functions of knowledge ascriptions. Identifying these implications will help me to make the case that extant versions of conversational contextualism should be rejected. In the next chapter I'll focus on the first of the specific versions of conversational contextualism that we'll look at; DeRose's gap view. I'll argue that DeRose's view

makes possible an objectionable form of trolling that undermines the truth of intuitively plausible knowledge claims. When real world power relations are thrown into the mix it's revealed that the ability to troll tracks asymmetries in power, which translates into asymmetries in who can deprive who of the important statuses and permissions that come with being ascribed knowledge that we've explored in this chapter.

Chapter 3: Scorekeeping Trolls

1. Introduction

I've said that removing the ideal assumption that speakers have equal abilities to influence the standards for 'knows' and replacing it with the more realistic non-ideal assumption that speakers have unequal abilities to influence the standards for 'knows' that track unfairly distributed power has three payoffs. It allows us to see unrecognized problems with extant versions of conversational contextualism, it helps us to identify an unrecognized form of epistemic injustice, and it motivates new contextualist theories. This chapter is the first of three that deliver on the first payoff. The focus is on a version of conversational contextualism that Keith DeRose (2004, 2009) has developed in response to Mark Richard's (2004) criticism that contextualism is unable to vindicate intuitions about disagreement. Richard's criticism is that if contextualism is true then two speakers who intuitively appear to be disagreeing with each other when one utters 'S knows that p' and the other utters 'S does not know that p' are not really disagreeing. If they have different standards in mind they are simply affirming and denying different propositions. DeRose responds by proposing the *gap view* (2009: 128-151) which is a meta-semantics for 'knows' that is designed to account for these intuitions about disagreement. According to this view, knowledge ascriptions are true iff the epistemic standards of each conversational participant are met, false iff each participant's standards aren't met, and truth-valueless otherwise. An implication of the gap view is that people with divergent standards can enter conversations and thereby render knowledge claims 'gappy' i.e. truth-valueless. I characterize this as a form of *trolling*. I argue that trolling results in unacceptably counterintuitive implications and that this constitutes a *reductio* against the gap view. Anyone can troll in principle, and so recognizing the problem of trolling does not depend on correcting for the

contextualists' neglect of power. But, paying attention to how power influences speakers' conversational abilities shows that it's not the case that anyone can troll in practice. There are asymmetries in who can troll that track asymmetries in power, and this translates into asymmetries in who can deprive who of the statuses and permissions that come with knowledge ascriptions that were outlined in the previous chapter. I'll also briefly explore the implications of trolling for other contextualist views about 'knows'.

In section 2 I describe the disagreement objection. In section 3 I describe DeRose's solution to the disagreement objection: the gap view. In section 4 I draw out counterintuitive implications of the gap view. In section 5, I show that the gap view, in conjunction with various knowledge norms, leads to further counterintuitive implications. I take these implications to constitute a *reductio* of the gap view. In section 6 I show how there are asymmetries in who can troll that problematically track power asymmetries. In section 7 I raise the concern that the gap view may not achieve its aim of vindicating intuitions about disagreement. Section 8 summarizes the chapter and previews what is to come in the next.

2. Contextualism and the Disagreement Objection

Recall from chapter 1 that contextualism is often motivated by soliciting intuitions about pairs of cases. DeRose is one contextualist who makes use of this strategy:

LOW: Hannah and Sarah are discussing whether the bank is open on Saturday. Hannah recalls 'I've been in there on a Saturday before'. She says, 'I know that the bank is open on Saturday'. Sarah agrees.

HIGH: Hannah and Sarah are discussing whether the bank is open on Saturday. Hannah says, ‘I’ve been in there on a Saturday before’. But Sarah replies ‘Banks can change their opening hours. You don’t know that the bank is open on Saturday’. Hannah agrees (variation on McKenna 2017, variation on DeRose 1992).

Assume Hannah is in the same epistemic position in both cases, and that the bank is open. Intuitively, DeRose claims, both Hannah’s knowledge ascription in **LOW** and Sarah’s knowledge denial in **HIGH** are true. Contextualists explain this by claiming that the truth-conditions of knowledge ascriptions are context sensitive because the epistemic standards that S must meet, in addition to having a true belief, in order for a knowledge ascription to be true depend on the conversational context in which the ascription is made. Hannah’s knowledge ascription is true given the epistemic standards that are operative in **LOW**. In **HIGH**, Sarah’s knowledge denial is also true, given the operative epistemic standards.¹² As we know from chapter 1, a common contextualist line is that speakers can change the operative epistemic standards by making alternatives to the proposition under discussion salient (e.g., Cohen 1988, DeRose 1995, Lewis 1996). When Sarah says ‘Banks do change their opening hours’ in **HIGH** she makes salient a possibility that must be ruled out if it’s to be truly said that Hannah ‘knows’ that the bank is open.

While contextualists take both Hannah’s knowledge ascription in **LOW** and Sarah’s knowledge denial in **HIGH** to be intuitively compatible, other cases involve knowledge ascriptions and denials that contextualists take to be intuitively incompatible. E.g.:

DISAGREEMENT: Hannah and Sarah are discussing whether the bank is open on Saturday. Hannah recalls it was open on a previous Saturday. She says, ‘I know that the

¹² Experimental studies have shown that intuitions about the bank cases that motivate contextualism aren’t uniform. See e.g. (Buckwalter 2017). In this chapter I am following DeRose’s presentation of contextualism. Nothing in my own argument in this chapter hinges on intuitions about the bank cases being as uniform as DeRose takes them to be.

bank is open on Saturday'. Sarah replies 'No you don't. Banks sometimes change their opening hours'. Hannah says, 'Oh come on, that's so unlikely'.

Intuitively, Hannah and Sarah are contradicting each other. Hannah is asserting the same proposition that Sarah is denying: they're disagreeing. The reason for intuitive compatibility in LOW and HIGH but incompatibility in DISAGREEMENT is that in contrast to LOW and HIGH, in DISAGREEMENT Hannah and Sarah are speaking to each other and are giving every indication that they're contradicting one another. Mark Richard (2004: 215-217) claims that the contextualist verdict on such cases must be that both assertions are true. Hannah's assertion is true given that her own personally indicated epistemic standards supply the content of her knowledge ascription, while Sarah's knowledge denial is also true given the content supplied by her own personally indicated epistemic standards. Given that, according to contextualism, Hannah and Sarah are expressing different contents in their respective uses of 'knows', contextualism delivers the counterintuitive verdict that there's no disagreement. Call this the disagreement objection.

3. The Gap View

DeRose accepts that a 'no disagreement' verdict on DISAGREEMENT would count against contextualism (2009: 130). DeRose claims that an effective response to the disagreement objection must meet two criteria:

- (1) Vindicate the intuition that there is disagreement in face-to-face interactions, and;
- (2) Ensure that each conversational participant's personally indicated standards register (2009: 145).

Meeting (1) is simply accommodating the intuition in cases like **DISAGREEMENT**. Meeting (2) is necessary for maintaining an essential feature of DeRose's contextualism. On DeRose's view, the conversational context determines the truth-conditions of knowledge ascriptions, which is partially determined by the personally indicated standards of conversational participants. So, a response that does not respect the personally indicated standards of all conversational participants would do away with the feature of contextualism that explains our intuitions about **LOW** and **HIGH**. This is precisely what motivates contextualism in the first place, and so would be a huge theoretical cost.

The structure of DeRose's response is familiar from supervaluationist theories of vagueness (e.g. Keefe 2008). In these theories, a vague predicate, e.g., 'bald', has several precisifications, 'has fewer than 100 hairs', 'has fewer than 200 hairs', etc. In evaluating the sentence 'Bob is bald', the sentence is first evaluated relative to each of the predicate's precisifications, at each of which the sentence comes out either true or false. The sentence is then supervaluated relative to this set of valuations, where it comes out as true simpliciter iff it's true at each valuation, and comes out as false simpliciter iff it's false at each valuation. Should it be true (false) at some valuations but not at others, then it's assigned no truth-value in the supervaluation. It's neither true nor false simpliciter.

On the gap view, precisifications are replaced by each conversational participant's individually indicated standards, and vague predicates are traded in for knowledge ascriptions. The truth-conditions for 'S knows the p' are as follows: 'S knows that p' is true iff the individual standards of every conversational participant are met, false iff the individual standards of every conversational participant are not met, and truth-valueless if some standards are met but others are not (2009: 144-145).

Consider the application of the gap view to **DISAGREEMENT**. Hannah and Sarah disagree about whether Hannah 'knows' that the bank is open on Saturday. On Sarah's standards, that banks can

change their opening hours is a salient error possibility, whereas for Hannah's it isn't. The sentence 'Hannah knows that the bank is open' is false relative to Sarah's standards and true according to Hannah's. Therefore, according to the conversational standard, it is neither true nor false. Contrast this with HIGH, where Hannah raises her standards in line with Sarah's. In HIGH, the sentence 'Hannah knows that the bank is open' is false relative to both speakers' standards, and therefore false simpliciter according to the conversational standard.

Let's evaluate DeRose's proposal against his two criteria. The gap view easily meets (2). Each of the conversational participants' personally indicated standards makes a difference in contributing to whether the proposition is true, false, or gappy. (1) also appears to be met. Because each participant is bound by the same conversational standard, the truth-conditions for 'S knows that p' are therefore the same when uttered by either, and thereby expresses the same propositional content. Each participant asserts or denies the same proposition and so count as disagreeing.

Given that the gap view meets both of DeRose's criteria, he takes himself to have proposed a successful response to the disagreement objection. For the remainder of this chapter, I draw attention to some unacceptably counterintuitive implications of DeRose's gap view. I argue that when taken together these implications constitute a *reductio* of the gap view.

4. Trolling

In this section I show that DeRose's gap view would, if correct, allow someone to enter a conversation with different epistemic standards from other conversational participants, and thereby undermine the truth of the other participants' knowledge claims. Call this phenomenon *trolling* and those who partake in it *scorekeeping trolls*. Scorekeeping trolls resemble their internet cousins

in that their presence in conversations derail the goals of their interlocutors. I show trolling to be a counterintuitive implication of the gap view that counts against it.

I'll illustrate trolling with two examples: first, one where the troll enters a conversation with higher epistemic standards than their interlocutors, and the second where they do the same but with lower standards.

VACCINE: While taking a riverside stroll with a friend, you're discussing their plans to take their children to be vaccinated against polio. Your friend remarks 'I know that vaccines reduce the chance of getting polio'. You agree. As you cross under a bridge, you encounter Anne te Vaks who asks what you're talking about. Your friend explains and repeats their claim. Anne denies that your friend 'knows'. She believes the reduction in rates of polio may be due to changes in diet and increased hygiene standards rather than vaccines. Her standards are much higher than yours and will only be met if your friend can refute these error possibilities. You do not defer to Anne's epistemic standards.

BRIDGE: Shane and Julia are driving through the mountains when they suddenly approach a narrow bridge. Beneath is a ravine with several burnt-out cars. As Julia starts to drive towards the bridge Shane shouts 'Stop! Look at those cars. You don't know that we won't fall off'. Julia takes her foot off the pedal, agreeing that she hasn't ruled out that possibility. While pondering what to do, they're approached by Alyssa, who asks them why they're not crossing. Shane explains, to which Alyssa replies 'Oh, don't worry, I've seen people cross before. You know you won't fall off'. Alyssa has much lower standards than Shane and Julia's, and they're met by her having seen cars successfully cross the bridge previously. Moreover, it's true that Shane and Julia would not fall off were they to cross, yet they don't defer to Alyssa's epistemic standards.

In both cases DeRose's gap view must say that the mere participation of the troll undermines ascriptions (denials) of knowledge, i.e., renders them gappy, due to the trolls' divergent personally indicated standards. Any further token utterances of the same knowledge ascription (denial) type will be gappy so long as the troll and the trolled party share a conversational context. This is a highly unintuitive result and will follow in any case with the following structure: a knowledge ascription is true (false) in a conversation because it meets (fails) each of the participants' epistemic standards. At a later time, a person enters the conversation with different standards, which are such that any future tokens of the same knowledge ascription type fail to meet (meets) these standards. This means that future tokens of the knowledge ascription type will be truth-valueless. This result occurs regardless of how outlandish or inappropriate the standards appear to be. The feature of DeRose's theory that makes trolling possible is that it meets his second criterion; each participant's standards must register, and this includes the troll's standards.

This counterintuitive implication ought to be of concern to DeRose. Throughout his work on contextualism it's clear that DeRose is in the business of capturing and explaining intuitions. DeRose takes one of the central motivations for contextualism to be its ability to explain our intuitions about cases such as **LOW** and **HIGH**. Furthermore, the motivation for proposing the gap view is precisely to vindicate our intuitions about cases of disagreement. It would be unprincipled and ad hoc for DeRose to be concerned to capture our intuitions in disagreement cases but to bite the bullet in trolling cases. Unless DeRose can explain away such counter intuitiveness then this should count against the gap view.

It might be objected that the trolls aren't conversational participants, and that therefore their utterances wouldn't affect the conversational score and undermine knowledge claims. While I won't settle on a theory of what it takes for someone to be a conversational participant here, I will

say something in favor of trolls' being conversational participants. We might plausibly say that a speaker is a conversational participant if their utterances register in conversational score (an informal register of moves made in a conversation). In the two cases just described, there is back and forth between the troll and trolled party, where the latter make no attempt to 'block' the troll's contribution. Rae Langton (2018) has argued that unblocked utterances are accommodated and thereby register in the score. Given this picture, the troll's utterances do register and thus they are participants.

5. Knowledge Norms

Further counterintuitive implications of DeRose's gap view can be derived when it's combined with various knowledge norms. Consider the following three contextualist formulations of much discussed knowledge norms that were introduced in the previous chapter:

ASSERTION (A): it is permissible for S to assert that p iff S satisfies 'knows that p' in S's context.

BELIEF (B): it is permissible for S to outright believe that p iff S satisfies 'knows that p' in S's context.

ACTION (PR): it is permissible for S to use p as a premise in her practical reasoning iff S satisfies 'knows that p' in S's context (Blome-Tillmann 2013).

Elsewhere, DeRose (2009: 108 – 109) endorses (A) and shows that it both provides a response to what he regards as a pressing objection to contextualism¹³ and a powerful argument in favour of

¹³ This is the warranted assertability manoeuvre objection to contextualism. See, e.g., Brown (2006).

contextualism (this is the argument that was described in chapter 2). Arguably, there is a parallel between assertion and belief such that the norms that apply to the former also apply to the latter (Williamson 2000: 255-256). Thus, if one accepts (A), as DeRose does, then one should also accept (B). While DeRose doesn't explicitly endorse (PR), he has gone to great lengths to demonstrate its compatibility with contextualism (2009: 247-251). For although (PR) is controversial, it has enough plausibility for its incompatibility with contextualism to count against the latter.

The interaction of these knowledge norms with the gap view compound the counterintuitive verdicts about our two trolling cases. In cases such as VACCINE and BRIDGE, your assertions, beliefs, and actions should be just as permissible after the troll has entered the conversation as beforehand. But, if the gap view is correct then assertions, beliefs, and actions that were previously permissible cease to be so once the troll enters the conversation.¹⁴ More than this, because knowledge claims are gappy in trolling contexts, the knowledge norms are no longer satisfied, and it's impermissible to assert, believe and act on the knowledge claims. To spell this out, notice that the knowledge norm of, e.g., assertion states that the permissibility of assertion depends on S satisfying 'knows that p'. S satisfies 'knows that p' iff 'S knows that p' is true. However, in trolling contexts 'S knows that p' is truth-valueless, i.e., not true. Therefore, S does not satisfy 'knows that p' and is not permitted to assert p.¹⁵ Consider VACCINE. According to (A) you are no longer licensed to assert 'Vaccines reduce the chance of getting polio', for 'You know vaccines reduce the

¹⁴ The claim here is not that what is, e.g., assertable, is wholly independent of one's conversational partners. Rather it's that the gap view, in conjunction with the knowledge norms, makes what's assertable too precarious.

¹⁵ In the current setup, a speaker S is only permitted to assert p iff 'S knows that p' receives the value true. In our examples, 'S knows (doesn't know) that p' does not receive a truth-value (it's truth-valueless) and hence isn't true. There may be alternative formulations of the truth-conditions for knowledge claims where the 'gappy' cases are such that it is indeterminate whether it is true (or false) rather than lacking a truth-value. In conjunction with the knowledge norms, these truth-conditions would imply that when one's 'knowledge' of p is 'gappy', i.e., indeterminate, the permissibility of assertion would itself become indeterminate. Although this may allow DeRose to escape my criticism, I will not explore it at length here. I note, however, that permissibility being indeterminate departs from the intuitive judgement that assertion etc. remain as determinately permissible after the troll enters the conversation as beforehand.

chance of getting polio' is no longer true. This is because the truth of the latter is required for the permissibility of asserting the former, according to (A). Next, applying (B), so long as Anne participates in the conversation, you ought not to believe, let alone assert, that vaccines reduce the chance of getting polio. Lastly, consider how (PR) applies in the example. Suppose you're on your way to vote on making polio vaccines mandatory, reasoning from your (previous) 'knowledge' that they reduce the chance of getting polio. (PR), in the presence of Anne, makes doing so impermissible. For, according to (PR), your 'knowledge' that vaccines reduce the chance of getting polio is required for using it as a premise in your practical reasoning.

In BRIDGE we have different consequences. There 'Julia doesn't know that she and Shane won't fall off' is rendered gappy. Assuming that prior to Alyssa entering the conversation, 'Shane and Julia know that Julia doesn't know that she and Shane won't fall off' is true, then this sentence is either rendered false or is itself gappy due to the gappiness of 'Julia doesn't know that she and Shane won't fall off'. The upshot of this is that Shane would no longer be permitted to assert that 'You don't know that we won't fall off'. The difference between these two cases is that whereas in VACCINE the high standards troll renders their interlocutors' 'knowledge' of some proposition p gappy, and therefore voids their license to assert, believe, and act upon p; in BRIDGE the low standards troll renders their interlocutor's 'knowledge' of the denial of the troll's 'knowledge' false or gappy, and therefore voids the interlocutor's license to assert etc. the denial.¹⁶

If DeRose's gap view is correct, there's reason to believe these consequences would be common. Self-ascribed standards for 'knowledge' often vary between people, and all the examples rely on one party entering a conversation and bringing with them higher (lower) standards than the other

¹⁶ I leave open whether when p is gappy, 'S knows that p' is false or gappy. In either case, asserting p becomes impermissible.

participants. On the contrary, I claim that the mere participation of someone in a conversation does not undermine the permissibility of one's assertions, practical reasoning, and beliefs. This would make the legitimacy of our actions, assertions, and beliefs too arbitrary. Someone might enter the conversation, thus voiding your license to act. Should they then leave, you would regain your license, only to lose it again when they return. Applied to VACCINE, suppose you always go by the bridge on the way to work and Anne always engages you in conversation about the same topic. It seems false that, in your context of ascription, 'You know that p', for a given p, becomes not true for a short period of time each day - whether 'You know that p' is true becomes fleeting, depending on who you chat to on the way to work. Apart from following from DeRose's gap view and the norms of knowledge, I see no reason to believe this.

Pairing the gap view with the knowledge norms amplifies the counterintuitive consequences of the gap view beyond those that were derived in section 3. Taken together, I claim, these consequences constitute a *reductio* of the gap view. Notice that any contextualist account of the truth-conditions of 'knows' that meets DeRose's second criterion - that each conversational participant's personally indicated standards register - will likely be susceptible to trolling. This is because meeting this criterion requires that the troll's standards register, thereby allowing for trolling.

6. Trolling and Power

Notice that in principle anyone can troll anyone. Because DeRose assumes that every conversational participant's personally indicated standards register, all it takes for a person to engage in trolling is for them to enter a conversation with different standards from existing participants, with the result that some standards are met, and others are not. This shows that recognizing the problem of trolling does not depend upon removing the idealized assumption that all speakers have

an equal ability to influence standards and replacing it with the more realistic non-ideal assumption that speakers have unequal abilities to shift standards that co-vary with unfairly distributed power, which is the central thought motivating this thesis.

But the reason why the significance of power has not shown up yet is because my critique has proceeded within DeRose's own framework, one that assumes an even linguistic playing field in which everyone's personally indicated standards register in determining the truth-conditions of knowledge claims. But this assumption is false. In practice, the linguistic playing field is not even, and everyone's personally indicated standards do not register. The powerful can often block the conversational moves of the powerless, preventing them from attaining the status of full conversational participant that is a requirement of trolling, where the reverse is not true. For example, in the context of a formal dinner at the University of Oxford the posh accented former Etonian Oxford student can successfully block the conversational move of the working class accented Northern student, but it's harder to imagine the opposite being true. During dinner conversation, the posh student declares 'I know that p!' The Northerner says, 'What about this not-p alternative?' To which the posh student replies, 'Maybe not-p seems plausible to people who went to a state school!' Not-p is not given uptake. The posh student hangs onto his knowledge claim. But suppose that it's the posh student raising the alternative to the working class Northern student's knowledge claim. The Northerner says, 'I know that p!' The posh student says, 'but what about not-p?' Not-p is given uptake. The Northerner's knowledge claim is undermined. What all of this means then is that those with power can troll, and those without will find it more difficult to do so. Asymmetries in who can troll mean that those with power can undermine the knowledge claims of the powerless, but the reverse is not true. Those with power can deprive those without of the statuses and permissions that come with being ascribed knowledge, and the reverse is not true. Where these asymmetries in power are unjust, so too will be the trolling that it facilitates.

7. Does the Gap View Dissolve Disagreement?

So far, my criticisms have focused on the consequences of the gap view following from it meeting DeRose's criterion (2). I've assumed that the gap view meets DeRose's criterion (1) by vindicating intuitions about disagreement. To finish up I'll cast doubt on whether the gap view does in fact vindicate these intuitions, thus failing to meet criterion (1). Consider the following:

WITHDRAWAL: Anya and Boris are epistemologists. They both hold the gap view and are feeling caffeine deprived. Boris says, 'I know that Keri Keri will have good coffee'. Anya denies this and raises an error possibility by uttering 'Perhaps only incompetent baristas have shifts today'. Boris rejects this as highly unlikely. Yet before the two continue disagreeing, they recognise that they are employing different epistemic standards. They take themselves as disagreeing, but as gap viewers they know that their respective knowledge ascription and denial are truth-valueless. In order to abide by the knowledge norm of assertion, both withdraw their knowledge claims.

Assume that: (i) relative to his own relatively lax standards Boris 'knows' he 'knows that Keri Keri will have good coffee'; (ii) relative to her own more demanding standards Anya 'knows' that 'Boris does not know that Keri Keri will have good coffee'; and (iii) that when first order 'knowledge' is gappy then second order 'knowledge' is either determinately lost or itself gappy. In either version of (iii), the assertion of first order 'knowledge' is impermissible. That is, when Boris's claim that 'I know that Keri Keri will have good coffee' is gappy then 'I know that I know that Keri Keri will have good coffee' is either false or is itself gappy, and thus Boris is not permitted to assert 'I know that Keri Keri has good coffee' because he does not 'know' this. When Anya's claim that 'Boris does not know that Keri Keri has good coffee' is gappy then 'I know that Boris does not know that Keri Keri has good coffee' is either false or is itself gappy, and thus Anya is not permitted to assert

'Boris does not know that Keri Keri has good coffee' because she does not 'know' this. What seems odd in WITHDRAWAL is that the disagreement has been dissolved only in virtue of the participant's awareness of and compliance with ASSERTION and the truth-conditions of knowledge ascriptions. The only way to avoid dissolving disagreement would be for them to violate these norms by not withdrawing their assertions. Compliance with the norms seems simply incompatible with being in a state of disagreement. Thus, the gap view appears to dissolve rather than vindicate disagreement, which was DeRose's aim in designing the view.

DeRose does comment on cases where conversational participants are aware of having different epistemic standards. His assessment is that in such cases each participant is bound by their own epistemic standards and are therefore both speaking truthfully when uttering their respective knowledge claims. In the cases that DeRose considers, the participants stipulate the different ways in which they are using 'knows' and for this reason do not take themselves to be disagreeing (2009: 132). Yet this analysis doesn't help to defend against the above criticism. First, the cases that form the basis of DeRose's analysis differ from WITHDRAWAL in that Anya and Boris, at first, do take themselves to be disagreeing and not merely stipulating different meanings of 'know'. Second, it is unclear why if you hold the gap view, you should favour the verdict where they are both speaking truthfully over the verdict that they utter gappy propositions. Importantly, whether it is mine or DeRose's assessment that is correct, disagreement is dissolved.

The gap view was meant to show how there could be genuine disagreement between conversational participants despite differing epistemic standards. However, what this objection shows is that there are cases where for the gap view to reach a verdict of disagreement the speakers need either to (a) be ignorant of the norms governing the practice or (b) intentionally violate them. It seems to me

that a successful response to the disagreement objection that is only successful conditional on pervasive ignorance and deviance isn't much of a success at all.

8. Summary and Preview

In this chapter I've argued that DeRose's gap view should be rejected because it allows for trolling, which results in unacceptably counterintuitive implications. Trolls can undermine the truth of intuitively plausible knowledge claims which, given the knowledge norms, is highly disabling as it allows trolls to undermine the permissibility of intuitively permissible assertions, beliefs, and actions. These problems are recognizable even before we take into account the effects of differences in power on abilities to influence the standards for 'knows'. But, once we do take into account the effects of power, the severity of the problems ratchet up even further. The powerful can troll more easily than the powerless, and thus can bring about the disabling effects of undermining people's knowledge claims such that these effects fall on the powerless. The ability of the powerful to troll the powerless contributes to and reinforces existing patterns of injustice and oppression. Thus, the critique is two-fold: (1) DeRose's gap view results in unacceptably counterintuitive implications and (2) when combined with real world power relations and the knowledge norms the gap view allows for a phenomenon that contributes to and reinforces injustice.

We thus have our first glimpse of how combining contextualism, real world power relations, and different assumptions about the function of knowledge ascriptions gives rise to epistemically, morally, and politically problematic implications. In the next chapter we'll see the next installment of problematic implications when we examine how Michael Blome-Tillmann's presuppositional epistemic contextualism fares when the ideal assumption that speakers have an equal ability to

influence standards is replaced with the non-ideal assumption that speakers have an unequal ability to shift standards.

Chapter 4: Power and Presuppositional Epistemic Contextualism

1. Introduction

The first payoff from combining contextualism, real world power relations, and different assumptions about the function of knowledge ascriptions is that it reveals an abundance of unrecognized problematic implications with extant versions of conversational contextualism. Chapter 3 provided the first installment of this first category of payoff by showing that Keith DeRose's (2004, 2009) gap view is vulnerable to an objectionable form of trolling. In this chapter I'll provide the second installment by showing that the success of Michael Blome-Tillmann's (2009, 2014) version of conversational contextualism, *presuppositional epistemic contextualism* (PEC), requires a level of co-operation among conversational participants that is not always forthcoming, and that even when it is, the terms of co-operation can be unfairly structured by unjust power relations. Recognizing this reveals a host of problematic implications of PEC that I'll argue are grounds for its rejection.

Whereas the gap view is a version of standards contextualism, PEC is a relevant alternatives version of contextualism. PEC says that which possibilities are relevant in determining the truth-conditions of a knowledge ascription depends on what is pragmatically presupposed in the context of ascription. I argue that the focus on pragmatic presuppositions is a mistake. Because what is pragmatically presupposed can be shaped by unjust power relations, PEC has a host of counterintuitive implications. The problematic implications are as follows: Firstly, I show that PEC is vulnerable to two classes of counterexamples that I label *authority problem cases* and *adaptive presupposition cases*. Secondly, I show that PEC, when combined with various plausible knowledge

norms, has a range of morally and politically concerning implications.¹⁷ Along the way I consider and reject some ways of developing PEC to deal with the problems I draw attention to.

By the end of this chapter, we will see a pattern starting to emerge: different authors respond to different problems for extant versions of conversational contextualism by shifting what features of conversational dynamics determine the content of ‘knows’ without placing adequate constraints on conversational dynamics. The upshot is that conversational contextualists need to either identify and impose adequate constraints on the effect of conversational dynamics on the content of ‘knows’, or divorce the content of ‘knows’ from conversational dynamics entirely. In chapters 7 and 8 I’ll explore the prospects of a version of each of these options. In chapter 7 I’ll explore the prospects of a virtue epistemology inspired conversational contextualism which constrains the effect of conversational dynamics on the content of ‘knows’ by only permitting those conversational

¹⁷ Jonathan Jenkins Ichikawa (2015) has developed a class of putative counterexamples to PEC. Ichikawa’s examples are ones in which PEC wrongly predicts that an intuitively false knowledge ascription, ‘S knows that p’, is true because S is ignorant of key propositions that are inconsistent with what is pragmatically presupposed in the context of ascription. For example: Emily says ‘The woman in the corner wants to speak to Sunil.’ Shari responds ‘Yes, she does. But, Sunil hasn’t looked at her in minutes, and didn’t seem to notice her then. He has no reason to even think she’s still there.’ Ichikawa claims that utterances of ‘Sunil knows that there is a woman in the corner’ would intuitively be false, but that given that it is pragmatically presupposed that there is a woman in the corner, error possibilities in which there is not a woman in the corner are incompatible with what’s pragmatically presupposed in the context, and are thus irrelevant according to PEC. PEC thus predicts that assertions of ‘Sunil knows that there is a woman in the corner’ are true (Ichikawa 2015: 1210). Ichikawa takes this to be a decisive reason to reject PEC.

Blome-Tillmann (2015) convincingly argues that Ichikawa’s counterexamples fail because while ‘knowledge’ undermining not-p possibilities are incompatible with what’s pragmatically presupposed, certain ‘knowledge’ undermining not-p possibilities are rendered relevant by other rules of relevance, in particular, the conjunction of the rule of actuality and the rule of resemblance (Lewis 1996). These not-p possibilities are not eliminated by S’s evidence, and so ‘S knows that p’ is false and PEC predicts the right result. This means that we are still on the lookout for successful counterexamples to PEC. My counterexamples fill this void.

Further, even if we assume that Ichikawa’s counterexamples are successful, we still learn something new about PEC from the counterexamples that I propose: whereas Ichikawa’s counterexamples show that PEC makes incorrect predictions because subjects are ignorant of key propositions, my counterexamples show that PEC gets the wrong result because it allows for the influence of power on what is pragmatically presupposed in conversational contexts. Additionally, the counterexamples I develop not only show that PEC makes inaccurate predictions, but also highlight politically problematic consequences of PEC.

moves that are underwritten by epistemically virtuous motivations to influence the truth-conditions of ‘knows’. In chapter 8 I’ll explore the prospects of a novel version of stakes contextualism that divorces the content of ‘knows’ from the conversational context altogether and instead ties it to what’s morally at stake in the context. I won’t endorse either of these views outright. My aim is more modest: to show that both avoid the problems that I identify for existing forms of conversational contextualism whilst retaining some of their benefits, and thus establish that both are worthy of further development.

The plan for this chapter is as follows: section 2 outlines PEC. Section 3 shows that PEC is vulnerable to two classes of counterexamples. I show that PEC has the resources to avoid the trolling problem that befell DeRose’s view, but that these very resources make it vulnerable to a different class of counterexamples: authority problem cases. I also develop a second class of counterexamples to PEC: adaptive presupposition cases. Section 4 shows that the counterexamples developed in Section 3 are not only epistemically problematic but are also morally and politically problematic. Section 5 concludes and foreshadows what is to come in chapter 5.

2. Presuppositional Epistemic Contextualism

Recall from the previous chapters that epistemic contextualism is the view that the truth-conditions of knowledge ascriptions vary with the context of the ascriber, and that conversational contextualism says that it is features of an ascriber’s conversational context that determines the truth-conditions of knowledge ascriptions. David Lewis (1996) proposed the following relevant alternatives version of conversational contextualism:

(L) 'S knows that p' is true in context C iff S's evidence eliminates every not-p alternative that is relevant in C.¹⁸

Lewis provided a series of 'rules of relevance' that specify which possibilities are relevant in C, where C is the context of ascription. If the rules determine that a possibility is relevant in a context, then it must be eliminated by S's evidence if S is to satisfy 'knows' in that context. On Lewis's view, it's these rules that determine how the content of 'knows' is influenced by contextual factors. The rule that concerns us is the rule of attention:

Rule of attention (RA): If an alternative is attended to in C, then it is relevant in C (1996: 569).

Lewis's commentary on (RA) is brief. He merely points out that (RA) amounts to nothing more than saying that an alternative to which attention is drawn is relevant (1996: 569).

Michael Williams (2001: 15) has pointed out that an obvious problem with (RA) is that it is too strong, in that it makes 'knows' too difficult to satisfy. To see this, consider the following case:

TEENAGER: Laurie is lying in bed. She hears her teenage son's bedroom window open, so she gets up and looks out. She sees someone who looks a lot like her son jumping out onto the ground below. She goes into the bedroom and finds it empty. She concludes that her son has slipped out, and spends the night fuming. In fact, this is exactly what happened. In the morning she challenges him: 'I know you went out late last night. I saw you.' The son responds: 'How do you know? It's possible that you dreamt the whole thing.'

¹⁸ The way in which Lewis originally words (L) is slightly different. It is formulated in terms of worlds being or not being properly ignored, rather than in terms of alternatives being or not being relevant. The substance of the two formulations is equivalent.

The dreaming possibility is not ignored, and so given (RA) it is relevant. What's more, the dreaming alternative cannot be eliminated by Laurie's evidence, and so Laurie's son is speaking truthfully when he denies that she 'knows'. This is a highly counterintuitive verdict; of course Laurie 'knows'! (RA) is thus too strong and should be rejected.

The problem that Williams identifies for Lewis's view has similarities to the trolling problem that I argued befell DeRose's view in the previous chapter. It's similar in that it involves someone entering a conversation, indicating that they have different epistemic standards from existing conversational participants, and thereby undermining the truth of existing participants' knowledge claims. But, the explanation for why this undermining happens differs for each of the views. It arises for DeRose's view because his proposed meta-semantics allows for trolls to render knowledge claims gappy. It arises for Lewis's view because (RA) allows conversational participants to raise farfetched alternatives and thereby render knowledge ascriptions false. Another difference is that whereas trolling, as a problem for DeRose, can be done both by those with higher and lower standards than existing conversational participants, the problem that Williams identifies for Lewis can only be caused by those with higher epistemic standards than existing conversational participants. It's a problem of undermining knowledge ascriptions only, and not knowledge denials.

Blome-Tillmann's view, PEC, is motivated by an attempt to avoid the problem that befell Lewis. PEC retains the core elements of Lewisian contextualism, but replaces (RA) with a 'rule of pragmatic presupposition' (2009: 249 - 56). In order to understand the rule of pragmatic presupposition, we need the notions of common ground and pragmatic presupposition.

Common ground: it is common ground that p in a group G iff all members of G accept (for the purposes of conversation) that p , all believe that all accept that p , and all believe that all believe that all accept that p , and so on (Stalnaker 2002).

Pragmatic presupposition: S pragmatically presupposes p in C iff S is disposed to behave, in her use of language, as if she believed p to be common ground in C (Blome-Tillmann 2009: 253, amended from Stalnaker 1974).¹⁹

Blome-Tillmann uses these two notions to formulate a replacement rule for (RA):

Rule of pragmatic presupposition (RPP): If an alternative is compatible with the speakers' pragmatic presuppositions in C, then it is relevant in C (2009: 248).

Pragmatic presuppositions are linguistic behavioral dispositions: they are dispositions to talk as if one had certain beliefs. This means that pragmatically presupposing p is compatible with not in fact believing that p. Because RPP has it that pragmatic presuppositions are dispositions to talk as if one has certain beliefs, speakers can voluntarily decide which propositions they presuppose. As soon as a speaker chooses to be disposed to utter a certain sentence with all that the sentence presupposes, it is brought about that they are in fact disposed to utter that sentence (2009: 254). Given that speakers can voluntarily decide what propositions they pragmatically presuppose, they thus have voluntary control over the content of 'knows' in a context (2009: 248). This means that PEC is well placed to deal with TEENAGER. So long as Laurie makes it clear that she won't consider the dreaming hypothesis, and her son acknowledges this and ceases to insist on it, the context will be one in which Laurie satisfies 'knows', even though attention has been drawn to an

¹⁹ Stalnaker's definition of pragmatic presupposition is cashed out in terms of belief: S pragmatically presupposes that p iff S believes p to be common ground (2002: 707 & 717). Because beliefs are not under voluntary control Stalnaker's definition of pragmatic presupposition leaves pragmatic presuppositions outside of the realm of the voluntary. This is a problem for two reasons. Firstly, as Stalnaker's own examples suggest, we can presuppose propositions even though we don't believe them to be common ground. E.g., suppose that I am in a foreign country in which I do not speak the language. I have no reason to think that the person that I approach on the street speaks English, but I am desperate, so I try: 'Is there a toilet nearby?' If I'm lucky then it will become common ground that we both speak English. My utterance presupposes that the person I approach speaks at least some English, even though I only hope that they do. If Stalnaker wants his definition of pragmatic presupposition to account for this case, then it requires amendment in the direction suggested by Blome-Tillmann (see Blome-Tillmann 2009: 249 - 252 for further discussion). Secondly, if Stalnaker's definition of presupposition is to be compatible with standard accounts of presupposition accommodation, then it must be amended in the direction suggested by Blome-Tillmann (see Blome-Tillmann 2009: 252 - 253 for discussion).

alternative. In this situation the dreaming hypothesis is incompatible with the pragmatic presuppositions of the conversational participants because its negation is mutually pragmatically presupposed. This means that the dreaming hypothesis is not relevant, and does not need to be ruled out for Laurie to continue satisfying ‘knows’.²⁰

3. Power and Pragmatic Presupposition

We’ve just seen that PEC appears to deal with TEENAGER. But notice that Laurie’s continuing to satisfy ‘knows’ depends on her son co-operating by ceasing to insist on the dreaming hypothesis. What if Laurie’s son doesn’t co-operate and continues to insist on the dreaming hypothesis? It is surely not uncommon for skeptics to sincerely insist on skeptical error possibilities even when it’s clear that their interlocutors won’t take such possibilities seriously. In these cases, Blome-Tillmann (2009: 268 - 272) claims that we end up in what Stalnaker (2002: 717 - 720) calls a ‘defective context’: a context in which the presuppositions of conversational participants diverge. In a defective context it is at best unclear which possibilities are relevant and at worst indeterminate. On the face of it, these features of PEC – the way in which it handles uncooperative interlocutors – would seem to make it vulnerable to a version of the trolling problem that was shown to befall DeRose’s gap view in the previous chapter. In this section I show that despite that initial appearance, PEC has the resources to avoid trolling, and so is able to handle the problems posed by uncooperative interlocutors in at least some cases. However, I show that the resources that help PEC to avoid

²⁰ It’s worth noting one further modification that Blome-Tillmann makes to Lewis’s view. Lewis famously does not require belief for the truth of knowledge claims. This is a widely rejected idiosyncrasy of Lewis’s view because it allows that ‘S knows that p’ can come out as true even when S’s belief is based on inappropriate epistemic grounds, such as being based on palm readings or on tea leaves. To avoid this implication Blome-Tillmann (2014: 31) adds that requirement that S’s belief must be properly based if ‘S knows that p’ is to be true.

trolling leave it vulnerable to two other classes of counterexample: authority problem cases and adaptative presupposition cases.

3.3. Trolling?

As we saw in the previous chapter, DeRose's (2004, 2009) gap view is vulnerable to an objectionable form of trolling whereby someone enters into a conversation with different epistemic standards from existing conversational participants, and thereby undermines the truth of their knowledge claims. The feature of DeRose's view that makes it vulnerable to trolling is that on his view the truth-conditions for knowledge ascriptions are determined by the personally indicated epistemic standards of those in the ascriber's conversational context. *Every* conversational participant's personally indicated standards register, no matter how outlandish or inappropriate those standards appear to be. PEC is structurally similar in that every conversational participant has a say in determining the content of 'knows', yet whereas what matters on DeRose's view is personally indicated epistemic standards, what matters on Blome-Tillmann's view are pragmatic presuppositions. *Every* conversational participant's pragmatic presuppositions contribute to the truth-conditions of knowledge ascriptions. It's easiest to see how this structural similarity makes PEC appear vulnerable to trolling once we amend the definition of trolling to reflect the difference between DeRose and Blome-Tillmann's views. Call the form of trolling that PEC seems susceptible to 'trolling PEC': PEC allows someone to walk into conversations with different pragmatic presuppositions from existing conversational participants and thereby undermine the truth of other participants' knowledge claims. To see the problem of trolling (PEC) more clearly, consider the following case:

SEXUAL HARASSMENT: Sarah has received an email from her boss, who is a man that she knows to have a history of committing sexual harassment, offering her a promotion. In the email Sarah's boss states that the promotion is conditional on her going out to dinner with him, after which they can go back to his hotel room and 'sign the contract'. Based on her knowledge of his history and the email she received, Sarah believes this to be a clear instance of quid pro quo sexual harassment, and her belief is true. Whilst in conversation with her friend, Sarah asserts 'I know that my boss sexually harassed me'. Steve overhears and joins the conversation. He knows of Sarah's boss's history of committing sexual harassment, yet he raises the possibility that on this occasion somebody hacked Sarah's boss's email account and sent the email to Sarah. This is not something that is ruled out by Sarah's evidence. Sarah and her friend both reject the possibility as being too outlandish to undermine Sarah's knowledge ascription, and so the possibility is clearly incompatible with their pragmatic presuppositions. Nevertheless, Steve insists on it.

It appears that PEC's verdict must be that it's unclear whether Sarah's knowledge ascription is true. This is because Steve's entrance into the conversation with stubborn pragmatic presuppositions that diverge from those of Sarah and her friend means that the context is defective. But, Sarah's knowledge ascription is clearly true. So, PEC gets the wrong result. Thus, PEC appears vulnerable to trolling PEC.

If, like DeRose's view, PEC is vulnerable to a form of trolling, then it ought to concern Blome-Tillmann. Fortunately, PEC has the resources to avoid the implication and so appears to be better than DeRose's view on this score. It's important to notice that in order for 'knows' to continue being satisfied in cases like **SEXUAL HARASSMENT**, the context must be one in which the negation of the skeptical hypothesis is mutually pragmatically presupposed. It's the fact that the

pragmatic presuppositions of the conversational participants in **SEXUAL HARASSMENT** diverge that means that PEC gets the wrong verdict on the case. Thus, if PEC is to avoid trolling, what's required is a strategy that ensures that when confronted with a skeptic, the context remains one in which the negation of the skeptical hypothesis is mutually pragmatically presupposed. Blome-Tillmann suggests one such strategy: an appeal to conversational authority (2009: 273). The appeal to the notion of conversational authority is a bid to ensure co-operation between conversational participants. Co-operation, in the form of mutually held pragmatic presupposition, is required if PEC is to capture intuitive verdicts on certain cases. But as we shall see as this chapter unfolds, speakers are not always co-operative, and even when they are the terms of co-operation can be unfairly structured by relations of domination and subordination which can itself be the source of problem cases.

Blome-Tillmann describes a conversational authority as someone who has 'control over the context in the sense that they have the power to make other participants to the conversation presuppose propositions' (2009: 273). In situations where there is a dispute as to whether some possibility is relevant in a context, a conversational authority can settle the issue by 'putting their foot down'. Blome-Tillmann explains that a conversational authority's attempt to put their foot down can have one of two outcomes. Firstly, if the audience is on the side of the conversational authority, then the skeptic will most likely be ignored by the group and be excluded from the context. The speakers thereby end up in a non-skeptical context, because the skeptic is no longer a conversational participant. Secondly, in a context in which there are participants with equal authority that clash, if none of the authorities are sidelined by the group, then we end up in a defective context (2009: 273).

Ending up in a defective context is what generates the problematic verdict in **SEXUAL HARASSMENT**, so the second aspect of Blome-Tillmann's conversational authority proposal is unhelpful. The first aspect of the proposal is more promising. While Blome-Tillmann doesn't spell out what it takes for someone to be a conversational authority, there are a variety of plausible ways that we might unpack the idea. One might possess conversational authority in virtue of possessing institutional authority. For example, an epistemology professor might put their foot down and insist that the negation of skeptical hypotheses be presupposed to stop students from getting bogged down in issues of skepticism. Alternatively, one might come to obtain authority by presupposing that one has it and having hearers go along with what one says. Rae Langton (2018) provides the following example to illustrate this form of authority:

There is a moment of indecision, and then someone takes charge; asks for suggestions about restaurants, decides on one, and asks someone to get two cabs while she calls to make reservations. When no one objects to this arrangement, she became the group leader, and obtained a certain authority. She did this by acting as if she had a certain authority (quoted in Langton 2018).

The attempt to obtain authority fails if others don't go along with it, and instead block the authority that was presupposed:

Someone could have objected, saying 'Who do you think you are, deciding where to go for us?' And the objection would have had a certain force (quoted in Langton 2018).

One might alternatively obtain authority by democratic vote. For example, one might propose excluding a particular speaker from a conversation and then asking, 'who's with me?' A particular person might be voted out of a conversation, forced to take their pragmatic presuppositions with them. An appeal to the last of these ways of attaining authority seems a promising way to deal with

Steve in **SEXUAL HARRASSMENT**. Sarah and her friend may decide that Steve is to be excluded from the conversation, and thereby ensure that that they remain in a context where the negation of the skeptical hypotheses is mutually pragmatically presupposed, and thus Sarah speaks truly when she asserts ‘I know that my boss sexually harassed me’. Thus, PEC has resources to avoid the counterintuitive verdict on **SEXUAL HARRASSMENT**.

3.2. Authority Problems

PEC appears to have the resources to escape trolling PEC; *merely* walking into a conversation with pragmatic presuppositions that diverge from those of other conversational participants is not sufficient to undermine the truth of existing participants’ knowledge claims.²¹ But, while the appeal to conversational authority helps PEC avoid problematic trolling in cases like **SEXUAL HARRASSMENT**, it is a source of problems in other cases. Consider the following:

SEXUAL HARRASSMENT*: Sarah has received an email from her boss, who is a man that she knows to have a history of committing sexual harassment, offering her a promotion. In the email Sarah’s boss states that the promotion is conditional on her going out to dinner with him, after which they can go back to his hotel room and ‘sign the contract’. Based on her knowledge of his history and the email she received, Sarah believes this to be a clear instance of quid pro quo sexual harassment, and her belief is true. During a company meeting, Sarah brings up the email and asserts ‘I know that my boss sexually harassed me’. While there are some possibilities that Sarah’s evidence cannot rule out, these are incompatible with her pragmatic presuppositions. Sarah’s boss is leading the meeting. He

²¹ Though, if trolls can be excluded by democratic vote, it’s possible for sceptical trolls to co-ordinate so as to outnumber conversational participants who are presupposing the negation of the sceptical hypothesis.

raises the possibility that somebody hacked his email account and sent the email to Sarah. This is not something that is ruled out by Sarah's evidence. Given the meeting leader's institutional position, he is the conversational authority. It is well known that the boss dislikes being challenged. Those who have done so in the past have faced professional repercussions. For this reason, the other meeting participants go along with what the boss says. He uses his conversational authority to settle the issue of what's to be pragmatically presupposed in the context, and the skeptical hypothesis is rendered relevant.

PEC's verdict in *SEXUAL HARASSMENT** must be that Sarah's knowledge ascription is false (and if the boss were to assert 'You do not know that I sexually harassed you', he would be speaking truthfully), which is the wrong result: it's clear that Sarah's knowledge ascription is true. What the example shows is that conversational authority can be wielded to insist on skeptical hypotheses, or keep us out of contexts where the negation of skeptical hypotheses are mutually pragmatically presupposed, and thereby be used to counterintuitively undermine the truth of knowledge ascriptions. Not only does PEC get the wrong verdict on the truth of Sarah's knowledge ascription, but the reason why PEC reaches this verdict is of the wrong kind to determine whether someone 'knows'. The reason why, contra intuition, Sarah's knowledge ascription is false is that someone powerful dictates what's to be presupposed in the context. But, the fact that someone is powerful is the wrong kind of reason for why someone should or should not count as 'knowing'. This unintuitive result will arise in any case with the following structure: intuitively 'S knows that p' is true, conversational authority is wielded to ensure that an ineliminable skeptical hypothesis is mutually pragmatically presupposed, and so PEC reaches the verdict that 'S knows that p' is false. The exclusion strategy that was helpful in dealing with *SEXUAL HARASSMENT* is not available here because the nature of the authority is such that the person cannot just be excluded from the conversation. Thus, we have our first class of counterexamples to PEC; *authority problem cases*.

One might wonder whether the meeting participants would ever really go along with what the boss says by shifting their pragmatic presuppositions in the direction he suggests. Even if the meeting participants don't directly call out the boss, they may make use of subtler strategies to communicate their non-acceptance of what he says. Situations are common in which a man makes a sexist remark, the women in the conversation say nothing, but they do each issue a roll of the eyes that's noticed by one another to be a signal that they notice and reject then man's sexist remark. This kind of strategy might be employed in situations like SEXUAL HARRASMENT*. The existence and use of eyeroll strategies show that problematic uses of power don't *always* succeed in changing pragmatic presuppositions.

But in our non-ideal world problematic uses of power can and do at least *sometimes* succeed in changing pragmatic presuppositions in the direction that the authority wants them to, and this is all that's needed to generate authority problem cases for PEC. In some circumstances, the other meeting participants would go along with what the boss says. Bosses are powerful people; they can help or hinder one's career, and they have the power to hire and fire. This can often be enough for people to go along with what they say, even if they vehemently disagree with the boss in the privacy of their own minds. An additional reason to think that it is plausible that other meeting participants would go along with what the boss says is that it's not uncommon for people to believe sexual harassment myths: that people lie about sexual harassment and that people see sexual harassment where there is none.²²

²² For example, one recent study published by the nonprofit Stop Street Harassment in conjunction with UC San Diego's Center on Gender Equity Found that 6% of people believe that most allegations of sexual assault or sexual harassment were lies for attention or money. This figure rises to 8% amongst men. Accessible at: <http://www.stopstreetharassment.org/wp-content/uploads/2012/08/2019-MeToo-National-Sexual-Harassment-and-Assault-Report.pdf>

Even if we assume that the other meeting participants do not go along with what the boss says, other ways that the conversation might go also spell trouble for PEC. Firstly, other conversational participants might not submit to the boss's authority and instead continue insisting on the negation of the skeptical hypotheses. In this situation the boss can wield authority to exclude dissenters from the conversation and ensure that it remains one in which the skeptical hypothesis is pragmatically presupposed. Secondly, we might imagine that there is an additional conversational participant who is of equal conversational authority to the boss. Blome-Tillmann (2009: 273 - 274) discusses conversations in which there are two equally authoritative conversational participants. One authority insists on a skeptical hypothesis, while the other refuses to take such a hypothesis seriously. Here, we end up in a defective context where it is unclear whether the knowledge claim is true. But, given that Sarah's knowledge ascription is true, this the wrong result. Any way we go with the case, conversational authority can be wielded in a way that generates a counterintuitive verdict.

What forms of power afford the conversational authority that is required for generating authority problem cases? The answer to this question will have implications for the range of cases that trouble PEC. A schematic answer is that it is power of the form and magnitude that gives one conversational authority that is equal to or greater than that of any other conversational participant. Institutional power such as that held by the boss in **SEXUAL HARASSMENT*** is a case that clearly fits this schema, and there are likely to be other forms of power fitting the schema too.

Is there a way to develop the notion of conversational authority that enables PEC to avoid problematic verdicts on authority problem cases? Let's consider three possibilities. Firstly, we might distinguish between legitimate and illegitimate exercises of conversational authority and claim that only legitimate exercises of conversational authority can make an alternative relevant. We might say

that legitimate exercises of conversational authority are attempts to settle what is pragmatically presupposed that are motivated by a desire to get at the truth, and that illegitimate exercises of conversational authority are attempts to settle what is pragmatically presupposed that are motivated by anything other than a desire to get at the truth. The boss's exercise of conversational authority is illegitimate, and so does not make the skeptical alternative he raises relevant. The problem with this strategy is that given that people can be wildly mistaken about what the truth is, their epistemic motivations can lead them to insist on presuppositions that generate problematic verdicts, such as Sarah's knowledge ascription being false. For example, we can imagine a variant of the case in which the conversational authority who insists on it being pragmatically presupposed that it's possible that the boss's email account was hacked is motivated by their belief that there's a conspiracy theory to get the boss criminally convicted of sexual harassment, or by their belief in ideologically motivated background claims that support rape culture, such as that women lie about being sexually harassed (see ft. 7).

Secondly, we could say that a conversational authority is an epistemic authority in Linda Zagzebski (2013) and Katherine Dormandy's (2018) sense. On this way of thinking, a subject is an epistemic authority for you when they are better than you at getting to the truth about some proposition because they have many more reasons pertaining to the issue than you and are more skilled at evaluating those reasons (Dormandy 2018: 773 - 774). Given her experiences, Sarah might be an epistemic authority on sexual harassment. Sarah may thus authoritatively declare that the skeptical alternative raised by her boss is irrelevant. Not needing to rule out the skeptical alternative, Sarah's knowledge ascription would be true, and PEC would deliver the right result. But, even if we suppose that Sarah is an epistemic authority, it is unclear why we should take someone being an epistemic authority in Zagzebski and Dormandy's sense to also be an authority on which alternatives are relevant. It simply doesn't follow from the fact that someone both has more reasons than me

bearing on the truth of *p* and is good at evaluating those reasons that they have authority over what alternatives *I* have to rule out to have a justified belief. Someone being an epistemic authority might put them in a strong position to determine whether relevant alternatives have been eliminated, but this does not put them in a position to determine which alternatives are relevant.

A third possibility is inspired by recent work on moral encroachment; the idea that moral considerations can affect the justificatory status of beliefs (see Bolinger 2020 for an overview). One way to think about moral encroachment is to think that the higher the moral stakes of someone forming a false belief, the more evidence that's needed for the belief to be justified (see e.g., Fritz 2017 for someone who takes moral encroachment in this direction). Another way that moral considerations might affect the justificatory status of belief is by keeping the evidential requirements for justified belief low when the moral stakes of failing to form a true belief are high. In chapter 8 I'll rely on this thought to develop a moralized version of epistemic contextualism that handles the problems levelled at extant versions of conversational contextualism in chapters 3, 4, and 5. For now we're just interested in whether this thought can be leveraged to develop the notion of conversational authority in a way that helps PEC avoid authority problem cases. We might say that a conversational authority is a moral authority, where a moral authority is someone who can reliably track the moral stakes of forming false beliefs. Now, if there is moral encroachment, moral authorities are plausibly well placed to see where moral considerations set evidential standards. Someone who knows that failing to form the true belief that Sarah suffered from sexual harassment is morally costly is well placed to see that evidential standards should be kept low. Armed with this insight, such a person can declare the skeptical alternative irrelevant in **SEXUAL HARASSMENT***. But what if there is no moral authority around to settle what's to be mutually pragmatically presupposed when presuppositions are misaligned? It seems that PEC must say that we find ourselves in a defective context, a result we're looking to avoid. Having examined three

ways of developing the notion of conversational authority and shown that each fall short, it's hard to see a principled way of developing the notion that would allow PEC to avoid authority problem cases.

3.3. Adaptive Pragmatic Presuppositions

We've just seen that authority problem cases present a class of counterexamples to PEC. In this section we will see that PEC is vulnerable to a second class of counterexamples that depend not on explicit exercises of conversational authority as in authority problem cases, but on the fact that people's pragmatic presuppositions are sometimes a function of unjust relations and distributions of power. What we see in these cases is that even when conversational participants co-operate, problem cases can still arise because the terms of co-operation can themselves be unfairly structured by relations of domination and subordination. Consider the following case:

SEXUAL HARASSMENT:** Sarah has received an email from her boss, who is a man that she knows to have a history of committing sexual harassment, offering her a promotion. In the email Sarah's boss states that the promotion is conditional on her going out to dinner with him, after which they can go back to his hotel room and 'sign the contract'. Based on her knowledge of his history and the email she received, Sarah believes this to be a clear instance of quid pro quo sexual harassment, and her belief is true. Whilst at a party, Sarah relays the content of the email to her friend. Sarah asserts, 'I know that my boss sexually harassed me'. Another party goer, Steve, overhears the conversation and joins in. Steve also knows of Sarah's boss's history of committing sexual harassment, yet he raises the possibility that on this occasion someone hacked Sarah's boss's email account and sent the email to Sarah. This is not something that is ruled out by Sarah's evidence. While Sarah believes

this possibility to be too far-fetched for it to be relevant, she's aware that Steve is a misogynist and fears how he will respond if she challenges him. Sarah does not challenge the relevance of the possibility that Steve raises, but would have if she were not concerned about Steve's response.

Recall that pragmatic presuppositions are linguistic behavioral dispositions, and as such can come apart from what an individual believes. Whilst Sarah believes that the possibility raised by Steve is too farfetched to be relevant, her linguistic behavioral dispositions are such that the skeptical hypothesis is compatible with her pragmatic presuppositions, and so compatible with the presuppositions of all conversational participants and thus relevant in the context. Given that the possibility is not ruled out by Sarah's evidence, PEC's verdict on SEXUAL HARASSMENT** is that Sarah's knowledge ascription is false. This is a problem for PEC for two reasons. Firstly, as in SEXAUL HARASSMENT*, it's clear that Sarah's knowledge ascription is true, and so PEC reaches the wrong verdict. And, secondly, as is the case with SEXUAL HARASSMENT*, PEC reaches the wrong verdict for a reason that is of the wrong kind. Whether or not Sarah 'knows' should in no way depend on her being scared of the responses she receives from those she's talking to. If PEC is true, then given that speakers are forced to pragmatically presuppose in line with the misogynist, they come out as not 'knowing'. This unintuitive result will arise in any cases with the following structure: intuitively 'S knows that p' is true, someone other than S successfully alters what is mutually pragmatically presupposed in a skeptical direction, this success depends on an unjust distribution of power, as a consequence PEC reaches the verdict that 'S knows that p' is false. Thus, we have our second class of counterexamples to PEC; *adaptive presupposition cases*.

It's reasonable to assume that cases like SEXUAL HARASSMENT** are common. As Crewe and Ichikawa (forthcoming) point out, we live in a rape culture - 'a cultural environment where sexual

assault and sexualized violence is a normative or expected type of interaction’ – which has harmful epistemological consequences. Claims about cases of sexual assault are frequently met with skepticism; ‘himpathy’ – the phenomenon in which sexual assault controversies foremost highlight the harms men and boys will experience as a result of being accused and found guilty of sexually assaulting someone – is widespread; and the threat of misogynistic policing mechanisms are ever present for those who threaten the patriarchal order (Manne 2017). It’s plausible that one way in which these forces operate is to shape the pragmatic presuppositions of sexual assault victims in a direction that makes ascriptions of knowledge of their own sexual assault come out as false. If you know that you’re likely to face skepticism, be made to listen to expressions of sympathy for your assailant, and risk hostile policing, you’ll be forgiven for taking steps to avoid such reactions. Such avoidances can involve pragmatically presupposing that the skeptical error possibility is relevant, and if PEC is correct it follows that ascriptions of knowledge to you of your own sexual assault are false.²³

4. PEC, Politics, and the Knowledge Norms

Beyond the epistemically counterintuitive verdicts that PEC reaches on **SEXUAL HARASSMENT*** and **SEXUAL HARASSMENT****, the two cases highlight politically troubling implications of PEC. Some political problems are implicit in the analysis of the cases given above: it’s plausibly an epistemic injustice that whether or not someone ‘knows’ depends upon illegitimate exercises of authority and unjust distributions of power. In this section I make explicit some

²³ We might also note that **SEXUAL HARASSMENT***, and real-world correlates of the case, which are surely pervasive, cast doubt on Blome-Tillmann’s assumption that pragmatic presuppositions are voluntary. If what one pragmatically presupposes is shaped by a social, historical, and political context that is characterized by unjust relations of power, we might be suspicious as to whether those pragmatic presuppositions are accurately characterised as voluntary. Voluntary or not, it doesn’t make a difference to the force of my counterexamples.

additional connected political problems with PEC. These problems are present in both **SEXUAL HARASSMENT*** (authority problem cases) and in **SEXUAL HARASSMENT**** (adaptive presupposition cases). I'll focus the discussion on **SEXUAL HARASSMENT***, but the problems might just as well have been illustrated by focusing on **SEXUAL HARASSMENT****.

At the core of the political problems facing PEC highlighted by **SEXUAL HARASSMENT*** is that conversational authority can be exercised in ways that undermine intuitively true knowledge ascriptions, which in turn deprives agents of statuses and permissions that they're entitled to. To see this, recall from chapter 2 that knowledge ascriptions have various socially valuable functions. Let's again consider the following three knowledge norms:

ASSERTION (A): it is permissible for S to assert that p iff S satisfies 'knows that p' in S's context.

BELIEF (B): it is permissible for S to outright believe that p iff S satisfies 'knows that p' in S's context.

ACTION (PR): it is permissible for S to use p as a premise in her practical reasoning iff S satisfies 'knows that p' in S's context (Blome-Tillmann 2013).

Consider how the knowledge norms will operate in **SEXUAL HARASSMENT***. According to (A) Sarah is no longer licensed to assert 'My boss sexually harassed me' in the context of the company meeting, for her utterance that 'I know that my boss sexually harassed me' is no longer true in that context. Note firstly that this is counterintuitive: given that it's clear that Sarah 'knows', it's clear that she's not violating (A) in asserting that her boss sexually harassed her. What's more, whether Sarah is permitted to assert that 'My boss sexually harassed me' should not depend on the presuppositions of her boss. Note secondly that this consequence is politically problematic. It

seems uncontroversial to say that it's in Sarah's interest to be able tell others about having been sexually harassed by her boss. Sharing this information can foster a sense of solidarity, it could give Sarah the confidence needed to take action and could provide others with the information needed to act on this fact too. But, if both PEC and (A) are true then Sarah cannot tell anyone about having been sexually harassed by her boss, in the context of SEXUAL HARASSMENT*, without opening herself up to criticism for violating (A). Not only does this diminish the likelihood of Sarah attaining the goods to be had from asserting the proposition, but it also imposes an additional burden of silencing on Sarah: Sarah cannot speak, lest she be criticized for violating norms of permissible assertion.

The implications of PEC and (A) suggest a novel form of silencing that's conditional on the conjunction of PEC and (A): it occurs when a conversational authority puts their foot down to insist on particular presuppositions in the context, which makes it the case that 'knows' is not satisfied, and given the knowledge norm of assertion, the person is no longer permitted to assert the proposition lest they open themselves up to criticism.

Now consider (B) in the context of SEXUAL HARASSMENT*. According to (B) Sarah ought not to believe that her boss sexually harassed her, let alone assert it. It's not even something that she can continue outright believing to herself in the context of the company meeting. Note firstly that this is counterintuitive: clearly Sarah is permitted to believe that her boss sexually harassed her, and whether she's permitted to believe that she was sexually harassed shouldn't depend on the presuppositions of her boss (who is her sexual harasser!). Note secondly that this consequence is politically problematic. If Sarah acknowledges that she is not permitted to believe that she was sexually harassed she could come to doubt herself. Mechanisms that undermine sexual harassment victims' beliefs in their own harassment are problematic. We might think of this as a novel form of

gaslighting (see, e.g., Abramson 2004, Stark 2019, Podosky forthcoming for philosophical discussion of gaslighting) that's conditional on the conjunction of PEC and (B). It occurs when a conversational authority puts their foot down to insist on particular presuppositions in the context, which makes it the case that 'knows' is not satisfied, and given the knowledge norm of belief, the subject of the knowledge ascription recognizes their own criticizability in continuing to believe the proposition and subsequently comes to doubt themselves.

Finally, consider (PR) in the context of SEXUAL HARASSMENT*. Suppose that Sarah is considering acting by making a formal complaint against her boss. Suppose also that the established administrative channel through which to make such complaints is in the company meeting. Given (PR), Sarah is permitted to act only on what she 'knows'. (PR) thus means that it is impermissible for Sarah to make a formal complaint against her boss given that she does not 'know' this to be the case in the context. Note firstly that this is counterintuitive: whether Sarah is permitted to act on the proposition that her boss sexually harassed her should not depend on the presuppositions of her boss. Note secondly that this consequence is politically problematic. Clearly sexual harassment is unjust and is something that we should want rid of. But, if the course of action that helps to get rid of it is p-dependent, and PEC generates the verdict that p is false, then this course of action cannot be uncritically pursued. The conjunction of PEC and (PR) thus makes sexual harassment hard to tackle.

Pairing the PEC problem cases with the knowledge norms compounds the problematic consequences that were highlighted in section 2. Not only does PEC get the wrong verdict about whether Sarah knows in SEXUAL HARASSMENT*, but it also has a host of politically problematic consequences when paired with the knowledge norms. Taken together, I submit, these consequences provide grounds for rejecting PEC.

One might try to take the sting out of the criticisms of PEC just presented by pointing out that Sarah only loses permission to assert, believe, and act on the proposition that her boss sexually harassed her in the context of SEXUAL HARASSMENT* and that once she moves to a new context, she regains her permissions. But this does little to dampen the concerns. As the above discussion shows, it is important that Sarah has the permissions in the context that she's in. The fact that she can regain the permissions in some other context is cold comfort.

A further objection might be that the politically problematic consequences just identified are not a consequence of PEC, but of PEC in conjunction with three controversial knowledge norms, and that if it can be shown that these norms aren't plausible then the defender of PEC should be unconcerned. There are two things to say here. Firstly, note that Blome-Tillmann (2013) has gone to great lengths to demonstrate the compatibility of PEC and the three knowledge norms I've considered. The formulations of the norms considered are in fact Blome-Tillmann's (2013) own. This gives us reason to think that he will be concerned with the consequences I've drawn attention to. Secondly, even if it turns out that knowledge ascriptions aren't connected to assertion, belief, and practical reasoning in the ways the knowledge norms suggest, there will be politically problematic consequences whatever the function of knowledge ascriptions. Knowledge ascriptions play socially valuable functions. There are various proposals of what these functions are: it has been suggested that knowledge ascriptions identify reliable informants (Craig 1990), or signal the appropriate end of enquiry (Kelp 2011), or track the epistemic norms governing assertion (Williamson 2000), belief, and practical reasoning (Hawthorne & Stanley 2008), or that knowledge ascriptions enable us to make important distinctions between blameless and blameworthy behavior (Beebe 2012). Whichever of these proposals turn out to be correct, people who find themselves in Sarah's position will not be able to do the socially valuable things with knowledge ascriptions that the proposals say are the function of knowledge ascriptions. This should count against PEC.

5. Summary and Preview

I've argued that PEC should be rejected because it is vulnerable to two classes of counterexamples - authority problem cases and adaptive presupposition cases - and because it has a host of morally and politically problematic implications. We're starting to see a pattern emerge. Both DeRose and Blome-Tillmann respond to problems for contextualism - DeRose to the disagreement objection and Blome-Tillmann to the problems faced by Lewis's rule of attention - by changing what features of conversational context the content of 'knows' is tied to - DeRose ties 'knows' to personally indicated standards and Blome-Tillmann to pragmatic presuppositions. Both philosophers, then, tie the content of 'knows' to features of actual conversations. But, both philosophers overlook the nasty real world nature of power as it operates in actual conversations and thus tacitly idealize the conversations that provide the data for their theories. As a result, both end up missing how the operation of real-world power relations affect the verdicts that their views deliver and both end up failing to place adequate constraints on how conversational dynamics influence the content of 'knows'. In the next chapter, I will move beyond critiquing specific versions of conversational contextualism to show that the failure to pay attention to how power operates in real world conversations leads to a further problem for all of the versions of conversational contextualism that we've looked at so far: conversational contextualism cannot diagnose clear cases of testimonial injustice.

In a recent paper, Robin McKenna (2017: 323) has gone further by claiming that the problem for conversational contextualism is not that it ties the content of 'knows' to the wrong features of conversational context, but rather that it ties it to conversational context whatsoever. McKenna (2017: 326 - 327) explains that conversational context is made up of the propositional attitudes of

conversational participants. Amongst these can be false propositional attitudes that function as defeaters of a subject's 'knowledge'. Thus, any version of contextualism that ties the content of 'knows' to the propositional attitudes of conversational participants, whether they be beliefs, pragmatic presuppositions, or the pragmatic presuppositions of particular speakers such as conversational authorities, is susceptible to counterexample. McKenna (2017: 329) tentatively suggests that conversational contextualism should be abandoned once and for all. I'm more optimistic. By removing the idealization that speakers have an equal ability to influence the content of 'knows' and replacing it with the non-ideal assumption that speakers have an unequal ability to influence the content of 'knows' that co-varies with unfairly distributed power, we've revealed unrecognized problems with extant versions of conversational contextualism. But, we get more than this. In revealing these problems we're also pointed in the direction of what an adequate constraint on the effect of conversational dynamics on the content of 'knows' might look like. All of the problem cases that we've looked at involve nefarious actors who raise alternatives for vicious reasons. Taking lead from this fact, in chapter 7 I'll explore the prospects of virtue epistemology inspired contextualism which constrains the effect of conversational dynamics on the content of 'knows' by only permitting those conversational moves that are underwritten by epistemically virtuous motivations to influence the truth-conditions of 'knows'. Thus there is greater room for optimism about conversational contextualism than McKenna would lead us to believe.

Chapter 5: Epistemic Contextualism or Epistemic Injustice?

1. Introduction

The previous two chapters identified problems for two specific versions of conversational contextualism: Keith DeRose's gap view, and Michael Blome-Tillmann's presuppositional epistemic contextualism (PEC). This chapter will identify a problem that faces all of the versions of conversational contextualism we've encountered. In a nutshell the problem is as follows: suppose that a subject, S, testifies to some 'known' proposition, p. Suppose further that the recipient of S's testimony, H, denies that 'S knows that p' because H is prejudiced against the social group to which S belongs. S has just suffered testimonial injustice at the hands of H. If there is a theory that provides those who harbor prejudice with the resources to avoid the charge of perpetuating testimonial injustice while still meeting speakers' testimony with a prejudicial response that generates testimonial injustice-like harms, then, regardless of what other virtues the theory has, alarm bells should be set ringing. I argue that conversational contextualism is such a theory.

Testimonial injustice is one form of the phenomenon of epistemic injustice which has received a great deal of attention in recent years - the idea that there is a distinctively epistemic class of injustices (see, e.g., Fricker 2007, Dotson 2012, Pohlhaus 2012, Medina 2013, Peet 2017). One way to understand testimonial injustice, and that way that I'm going to propose that we understand it in this chapter, is as involving wrongful knowledge denials. For a much longer period, contextualists have influentially argued that the truth-conditions of knowledge ascriptions - sentences of the form 'S knows that p' - are context-sensitive (see, e.g., DeRose 1992, Lewis 1996, Cohen 1999, Blome-Tillmann 2014, Pynn 2015, Ichikawa 2017). Contextualists claim that whether 'S knows that p' is true depends in part upon the epistemic standards that are put in place by the

conversational moves made by speakers in the ascriber's context. If particularly demanding epistemic standards are put in place, 'S knows that p' will often come out false. In such circumstances, whatever intuition tell us, if H denies that 'S knows that p', S cannot be said to suffer testimonial injustice. Testimonial injustice as we'll be understanding it in this chapter involves a wrongful denial of 'knowledge', a fate that cannot befall non-'knowers'.

These features of contextualism mean that if it is correct, many cases that are taken to be instances of testimonial injustice are in fact not. Contextualists are thus faced with a dilemma: either reject current versions of conversational contextualism in order to vindicate the intuition that such cases are in fact instances of testimonial injustice, or, hand over the resources that allow the prejudiced to avoid the charge of having perpetrated testimonial injustice and concede that such cases are not in fact instances of testimonial injustice. I embrace the first horn of the dilemma, but I do not abandon conversational contextualism entirely. Rather, I take the problem identified for current versions of conversational contextualism in this chapter (alongside the problems identified in the previous two chapters) to be part of a motivation for developing in a new form of conversational contextualism in chapter that avoids the problems.

This chapter brings the first payoff to be had from theorizing contextualism non-ideally to an end. By the end of this chapter, we will see that there is a general problem with the way that theories of conversational contextualism have been developed that explains the more specific problems that I've been identifying: existing versions of conversational contextualism fail to place adequate constraints on how conversational dynamics influence the truth-conditions of knowledge ascriptions. The upshot is that an acceptable contextualist account of knowledge ascriptions must either figure out a way to appropriately constrain the effect of conversational dynamics of the truth-conditions of knowledge ascriptions, or divorce the truth-conditions of knowledge ascriptions from

conversational dynamics entirely and tie them to some other feature of context besides conversational dynamics. I'll go on to explore the former option in chapter 7, and the latter option in chapter 8.

This chapter proceeds as follows: in section 2 I describe the phenomenon of testimonial injustice, propose a particular understanding of testimonial injustice, and provide an example. In section 3 I'll carry out a contextualist assessment of the example of testimonial injustice introduced in section 2. This assessment will give rise to the dilemma. In section 4 I consider three ways in which contextualists might attempt to resist the dilemma: (i) by rejecting my proposed understanding of testimonial injustice; (ii) by insisting that contextualism and testimonial injustice have fundamentally different subject matters, and; (iii) by demonstrating that the versions of conversational contextualism that we've been considering are sophisticated enough to avoid the dilemma. In section 5 I explain the general problem with conversational contextualism that underlies the more specific problems that have been identified across chapters 3, 4, and 5. This sets us up for developing new versions of contextualism in chapters 7 and 8.

Before we get started let me reiterate that the problem that I identify in this chapter is a problem for all of the versions of conversational contextualism that we've encountered. Thus, I'll be targeting both standards contextualism, in the form of DeRose's gap view, and relevant alternatives contextualism, in the form of Lewis's contextualism and Blome-Tillmann's PEC, and others besides that I haven't spoken about in as much detail (see references in chapter 1). As much as is possible, I'll talk in terms of standards when addressing standards contextualists and in terms of alternatives when addressing relevant alternatives contextualists. There will be some points at which it's difficult to avoid moving between the two models. Make use of the epistemic standards to relevant alternatives translation offered in chapter 1 at these points.

2. Testimonial Injustice

Testimonial injustice is one form of the phenomenon that Miranda Fricker has labelled epistemic injustice: a class of injustices that are distinctively epistemic in that the wrong consists in ‘a wrong done to someone specifically in their capacity as a knower’ (2007: 1). Testimonial injustice occurs when the testimony of a speaker receives an unfair deficit of credibility from a hearer owing to an identity prejudice - a prejudice against an individual in virtue of their membership in a particular social group - on the part the hearer (2007: 27). This, I think, is the way that most people have come to understand the phenomenon of testimonial injustice. We can define it as follows:

TI-Credibility: S suffers testimonial injustice at the hands of H if S testifies that p to H and H gives S less credibility than S deserves because H has an identity prejudice towards a social group to which S belongs.

One way that H can give S less credibility than S deserves is by wrongfully denying that S ‘knows’:

TI-‘Knowledge’: S suffers testimonial injustice at the hands of H if ‘S knows that p’ is true, S testifies that p to H, and H denies that ‘S knows that p’ because H has an identity prejudice towards a social group to which S belongs.

Notice that TI-‘Knowledge’ is defined in terms of knowledge ascriptions and denials, and not in terms of epistemic properties such as knowledge itself. As such, TI-‘Knowledge’ takes testimonial injustice to be a primarily linguistic or conversational phenomenon. This may raise some eyebrows. It might be objected that testimonial injustice concerns knowledge itself, not knowledge ascriptions and denials. The objection continues, testimonial injustice occurs when a hearer prevents a speaker from imparting the property of knowledge (or some degree of justification or warrant), not when a hearer wrongly takes a speaker to fail to satisfy the predicate ‘knows’! But, there are good reasons

to think that TI-‘Knowledge’ is, at the very least, *a way* that that testimonial injustice can be perpetrated. Rejections of testimony because of identity prejudice are paradigm cases of testimonial injustice and responding to a known assertion with a knowledge denial is one common way to reject testimony. What’s more, many of the harms of testimonial injustice are a result of the conversational moves made by hearers in response to speakers’ testimony. For example, the primary harm of testimonial injustice is a particular kind of insult (more on the details shortly). But speakers can only feel insulted if they have access to hearers’ assessments of their testimony, and speakers only have such access through observing the conversational moves made by their interlocutors.

All I’ve said so far is that TI-‘Knowledge’ is *a form* that testimonial injustice can take. In section 4 I’ll argue that some contextualists must accept that TI-‘Knowledge’ is *the way* to understand testimonial injustice. Before we get to that, I’ll set out this chapter’s main argument. When testimonial injustice is understood as TI-‘Knowledge’, both testimonial injustice and contextualism have the same subject matter: knowledge ascriptions and denials. Assuming the TI-‘Knowledge’ understanding of testimonial injustice, I’m going to argue that contextualism is unable to vindicate our intuitions about some paradigm cases of testimonial injustice.

Consider the following case of TI-‘Knowledge’. Call it **DISMISSED FROM CLASS**. A black public high school student named Rashaan has been sent to the Assistant Principal’s office by his teacher, Ms. H. Ms. H’s decision to send Rashaan out of her class was guided by racial prejudice.²⁴

²⁴ It is reasonable to suppose that this is true. It is well established that there are racial disparities in school discipline in the United States. Black students are 3 times more likely than white students to be suspended or expelled (Fenning & Rose 2007; Skiba et al. 2011; Skiba, Michael, Nardo, & Peterson 2002). Okonofua and Eberhardt (2015) have recently revealed the psychological mechanisms involved in bringing about this disparity. The studies that they have carried out show that teachers are likely to interpret students’ behavior differently based on race: teachers are more likely to perceive minor infractions by black students as being more severe than those of white students; report being more irritated by infractions of black students rather than white students; think it appropriate to punish black students more harshly than white students for equivalent infractions, and; are more likely to view black students, rather than white students, as troublemakers.

Suppose that Rashaan has a justified true belief, and so ‘knows’, that Ms. H’s decision to send him out was guided by racial prejudice. Rashaan has ample evidence for the proposition: he has witnessed other non-black students behave just as he did in Ms. H’s class and yet not get sent out. From these observations Rashaan infers that Ms. H’s decision to send him out was guided by racial prejudice. In addition to his observational evidence, we might add that after several years of repeatedly being the target of racial stereotyping and being sent to the Assistant Principal’s office for minor offences Rashaan has become attuned to patterns of racial discrimination in his school life. Numerous philosophers have persuasively argued that particular social locations often come with epistemic advantages. For example, when it comes to accessing knowledge about the racism faced by black people, black people are best positioned know about how and when it manifests, because they experience it frequently in a patterned way (e.g., Wylie 2004, Alcoff 2005, Mills 2007). If this is right, then this gives us additional reason to believe Rashaan when he testifies about his experience of racial discrimination.

Now suppose that the following conversation takes place when Rashaan reaches the Assistant Principal’s office:

AP: Rashaan, why have you been sent to my office?

R: I was trying to involve myself in the discussion, but Ms. H. would not give me a chance and sent me out. She never sends out white kids for trying to participate. This is racism!

AP: Rashaan, it’s possible you are mistaken. Can you rule out the possibility that you’ve misinterpreted Ms. H’s actions? It’s possible that you’re systematically misinterpreting the actions of others when directed at you!

R: Of course, it’s a possibility! But I’ve seen white kids not get sent out for doing the same

thing!

AP: If you can't rule out the possibility that you've misinterpreted Ms. H's actions, then you don't know that Ms. H's actions were guided by racism.

With Rashaan unable to eliminate the possibility that he has misinterpreted Ms. H's actions – this is not a possibility that can be ruled out by one's own observational evidence – the Assistant Principal judges Rashaan to not 'know' that Ms. H's decision to send him out was guided by racial prejudice. Assume that the Assistant Principal's judgement is guided by racial prejudice; the possibility that black people misperceive the actions of others as involving racist treatment is one that is all too commonly believed by those who harbor racial prejudice. The Assistant Principal's refusal to ascribe knowledge to Rashaan is a clear case of TI-'Knowledge' and of the harms that result when it is perpetrated.

On the basis of observing Ms. H. treat black and non-black students differently when they behave in similar ways, Rashaan forms the justified true belief, and so 'knows that Ms. H's decision to send him out was guided by racial prejudice', yet the Assistant Principal does not treat Rashaan as 'knowing' this to be the case; Rashaan's is not ascribed knowledge when it should be because of an identity prejudice held by the Assistant Principal. It does not require a great deal of imagination to see that the harms of testimonial injustice that Fricker identifies are likely to be present in **DISMISSED FROM CLASS:** Rashaan is prevented from imparting knowledge to the Assistant Principal. Fricker claims that because the capacity to impart knowledge is a distinctive human capacity, when one is undermined in this capacity one is disrespected qua human. Fricker calls this the primary harm of testimonial injustice.

Rashaan is also likely to suffer from what Fricker identifies as secondary harms of testimonial injustice: practical or further epistemic consequences of the primary harm. The practical harms

may manifest in a bad relationship taking hold between Rashaan and the teachers at his school, leading Rashaan to struggle to trust his teachers. This will have an adverse impact on his education, his ability to do well at school, and consequent ramifications throughout his entire life insofar as doing well at school matters for how well one's life goes. An example of a secondary epistemic harm is that when a hearer is not believed they may lose confidence in their belief, and so cease to satisfy a condition on knowledge (2007: 43 - 58).

3. Contextualism and Testimonial Injustice

In this section I will carry out a contextualist assessment of **DISMISSED FROM CLASS**. In so doing I will show that if contextualism has it right, then the truth-conditions of knowledge ascriptions can shift such that there are some contexts in which not attributing knowledge to Rashaan does not constitute testimonial injustice because Rashaan cannot be truly said to 'know that Ms. H's decision to send him out was guided by racial prejudice'. To see how this is so, let us first remind ourselves of contextualism and then carry out a contextualist assessment of **DISMISSED FROM CLASS**.

Contextualism is the view that the truth-conditions of knowledge ascriptions (and denials) vary according to the context of the ascriber. The truth-conditions of knowledge ascriptions are constituted by the epistemic standard that S must meet, in addition to having a true belief, to make the knowledge ascription true. Hence, what varies with context is the epistemic standard that S must meet in order for the knowledge ascription to be true. Accordingly, there are some contexts in which 'S knows that p' requires for its truth that S have a true belief that p and meet low epistemic standards, and there are other contexts in which the truth of the same knowledge ascribing sentence requires for its truth that S have a true belief that p and meet very high epistemic standards.

According to conversational contextualism, it is the conversational context in which the knowledge ascribing sentence is uttered that determines the operative epistemic standards. The conversational context is constituted by speaker intentions, listener expectations, presuppositions of conversational participants and salience relations; this is what David Lewis has called the ‘conversational score’ (1979). Changes in the conversational context induce a change in the epistemic standards that must be met if knowledge ascribing sentences are to be true. We’ve seen that there is some disagreement amongst contextualists about exactly how and when conversational contexts change to raise epistemic standards. I will return to this disagreement in section 4, but for present purposes we can assume the common contextualist view that epistemic standards are raised when moves in the conversation make salient various error possibilities that are alternatives to the proposition under discussion (Lewis 1996, DeRose 1992, Cohen 1999). If the skeptic makes salient the possibility that we are handless brains in vats, then we can no longer truthfully utter ‘I know I have hands’. This is because the standards that are now operative require us to eliminate the possibility that we are handless brains in vats if we are to count as ‘knowing that we have hands’, which is not something that we can rule out.

If contextualism is correct about the mechanism by which changes in context come about, then the conversational moves that are made by the Assistant Principal – raising the error possibility that Rashaan has misinterpreted Ms. H’s actions as part of a pattern of systematic misinterpretation of other people’s actions directed at him – shifts the truth-conditions for ‘Rashaan knows that Ms. H’s decision to send him out was guided by racial prejudice’. When this error possibility is introduced, the epistemic standards are cranked up just as they are in the case where the skeptic introduces the handless brain in a vat possibility. While Rashaan has the evidence that he does, he cannot eliminate the possibility that he has misinterpreted Ms. H’s actions. One cannot eliminate the possibility of having misinterpreted one’s own observations using those very observations. If Rashaan cannot

eliminate this possibility, then according to contextualism, it is not true to say that ‘Rashaan knows that Ms. H’s decision to send him out was guided by racial prejudice’. The result that contextualism yields conflicts with what we intuitively want to say about the case, namely, that Rashaan ‘knows’ and that denying this because of an identity prejudice constitutes testimonial injustice in the form of TI-‘Knowledge’.

More than just reaching the wrong verdict about **DISMISSED FROM CLASS**, contextualism can be shown to have provided the Assistant Principal with the resources to strategically avoid charges of perpetrating testimonial injustice. We can see this by filling out **DISMISSED FROM CLASS** as follows. Suppose that in his days as a student the Assistant Principal took philosophy courses in which he learned about both contextualism and epistemic injustice. Aware of his racial prejudices, the Assistant Principal is determined not to risk getting the sack from his job by perpetrating testimonial injustice against students. Luckily for him, the Assistant Principal recognizes that contextualism provides just the ticket. All he needs to do in order to avoid perpetrating TI-‘Knowledge’ is to raise error possibilities that induce a high standards context, prior to his passing judgement on the credibility of any student’s testimony. Once in the high standards context the Assistant Principal can safely deny that his interlocutors ‘know’ that which they testify to, free from the risk of being accused of perpetrating testimonial injustice. This is precisely the strategy employed in **DISMISSED FROM CLASS**. Contextualism implies that if hearers employ this strategy, they avoid testimonial injustice.

I want to turn now to a contextualist assessment of a second case of testimonial injustice in order to assuage any worries one might have that the implications of contextualism that I’ve drawn attention to only apply to isolated or merely fictional cases. This second case is a popular illustration in the testimonial injustice literature taken from Kristie Dotson’s work (2012). Dotson’s example draws

on a case from Patricia Williams' book *The Alchemy of Race and Rights* (1991). Williams describes an experience that she had of trying to enter a luxurious New York clothing shop. Williams tried to enter the shop in middle of the day during the Christmas shopping season. In order to gain entry shoppers had to press a buzzer before being let in. After pressing the buzzer Williams was denied entry. Williams reports that there were numerous white shoppers inside and that there was no other identifiable reason for not being allowed in besides racial prejudice; as an African-American she represented undesirable patronage. Dotson identifies the moment at which Williams suffers from testimonial injustice as coming later, when she tried to relay the incident to a class that she was teaching at Stanford Law School. A rumor got started amongst the class that her story was 'a lie that was probably the product of a diseased mind trying to make all of the white people feel guilty' (1991: 242). Williams's class were skeptical that the shop employees would deny someone entry because of their race. The students reached for alternative explanations, such as, 'perhaps all of the people inside worked there and were just on a break when you buzzed'. While Williams had the evidence that she did, she was not able to eliminate the error possibility raised by the student. Even if Williams were to point out that she could distinguish staff from shoppers by the all black uniform worn by the staff, the student might reply by pointing out that it is possible that there is no official uniform and that it might be a mere coincidence that multiple people were wearing similar black outfits. All black is a popular look after all. Once Williams's interlocutors begin raising hard to eliminate error possibilities, contextualism reaches the verdict that Williams doesn't satisfy 'knows'. Thus, the refusal to ascribe 'knowledge' to Williams does not amount to testimonial injustice. The contextualist verdict is at odds with that which has been reached by those in the testimonial injustice literature.

It might be objected that there is no special problem here, and that all I've done is reiterate the oft made charge that contextualism is too skeptic friendly. For example, suppose that I'm confident

that I saw my partner hanging out with my arch-nemesis last week. I report this to my friend, who says ‘are you certain that your arch-nemesis doesn’t have an identical twin that your partner is friends with?’ To which I reply, ‘Well, no, I guess I’m not’. I am unable to rule out the possibility, and so the contextualist verdict is that I don’t ‘know’ in this case. But intuitively, this is just too far-fetched a possibility to count against my ‘knowing that I saw my partner with my arch-nemesis’. The structure of the case is the same as that of **DISMISSED FROM CLASS**. There is a case in which intuitively we think that ‘S knows that p’ is true, but given the contextualist account of the mechanism by which epistemic standards shift, contextualists must say that ‘S knows that p’ is false. This problem was described in the previous chapter when discussing the failings of Lewis’s contextualism.

It would be a mistake to think that the inability of contextualism to account for intuitions about cases of testimonial injustice is merely an iteration of the charge that contextualism is too skeptic friendly. We can see this by distinguishing between two versions of the ‘too skeptic friendly’ charge. The first version is that which was just described in the previous paragraph; that as soon as a skeptical error possibility is introduced epistemic standards are dramatically raised such that one can no longer truthfully claim to ‘know’. Thus, contextualism is too skeptic friendly. Varieties of this objection have been made by, for example, Schiffer (1996), Barke (2004) and, Brendel (2005). The second version, and the one that I’m interested in here, is not merely a matter of contextualism being too friendly to the epistemological skeptic, but also involves a troubling additional skeptical implication about whether an injustice has taken place. Contextualism offers prejudiced recipients of testimony the resources to pull off the trick of avoiding being charged with perpetrating testimonial injustice while still meeting a speaker’s testimony with a prejudicial response that generates testimonial injustice like harms. I suspect that many will take this second version of the ‘too skeptic friendly’ charge against contextualism to be more pressing. While a theory that lets

epistemological skeptics get their way with ease might be thought to be worrisome, a theory that arms racists and sexists with the tools to get away with perpetrating race and sex-based injustices is likely to be thoroughly alarming!

We have seen that a contextualist assessment of **DISMISSED FROM CLASS** delivers the verdict that Rashaan ceases to suffer from testimonial injustice once particular error possibilities are made salient. This is because if contextualism is correct then it is not true to say that ‘Rashaan knows that Ms. H’s decision to send him out was guided by racial prejudice’ in such circumstances. Thus, conversational contextualists are presented with a dilemma: either they must reject current versions of conversational contextualism in order to vindicate the intuition that such cases are in fact examples of testimonial injustice, or, continue to endorse contextualism and concede that such cases are not in fact examples of testimonial injustice once epistemic standards are raised.

4. Resisting the Dilemma

Let’s now turn to consider three ways in which the contextualist might attempt to resist the dilemma that I have set up: the first stems from a rejection of my characterization of testimonial injustice as TI-‘Knowledge’; the second stems from an insistence on contextualism and testimonial injustice having fundamentally different subject matters; and the third involves attempting to demonstrate that the versions of conversational contextualism that are proposed by DeRose and Blome-Tillmann have the resources to avoid the dilemma.

4.1. Credibility Ascriptions, Knowledge Ascriptions, and Testimonial Injustice

The contextualist might attempt to resist the dilemma by objecting that TI-‘Knowledge’ is a

mischaracterization of testimonial injustice and to claim that a more accurate characterization does not have the implication that I have drawn attention to. The objection is as follows: testimonial injustice occurs when a hearer does not ascribe the testimony of a speaker the credibility that it is due, and this does not amount to a hearer not ascribing 'knowledge' to a speaker. A hearer can ascribe, or fail to ascribe, a speaker's testimony the credibility that it is due, regardless of whether they ascribe them 'knowledge'. So, even if it is an implication of contextualism that Rashaan ceases to 'know' once particular error possibilities are made salient, it does not follow that he does not suffer from testimonial injustice because it can still be the case that the Assistant Principal does not give Rashaan's testimony the credibility that it deserves.

Settling whether this response gets the contextualist out of trouble requires taking a closer look at what credible testimony consists in and what is involved when a hearer assesses the credibility of a speaker's testimony. The function of a credibility assessment is to determine how a hearer should adjust her beliefs when confronted with incoming testimony. Fricker claims that there are two considerations that a hearer, *H*, must bear in mind when assessing the credibility of a speaker, *S*'s, testimony that *p*: firstly, *H* must assess *S*'s sincerity with respect to *p*, and; secondly, *H* must assess *S*'s competence with respect to *p*. *H* judges *S* to be credible with respect to *p* just in case *H* judges *S* to be both sincere and competent with respect to *p*. *H*'s assessment of *S*'s testimony that *p* as credible is accurate just in case *S* is in fact both sincere and competent with respect to *p* (Fricker 2007: 32).

It is clear what it takes for *S* to be sincere when testifying to *p*; *S* must not be attempting to deceive *H*. Fricker is less clear about what it takes for *S* to be competent when testifying to *p*. At some points she writes that testimonial injustice occurs when a speaker is undermined as 'a giver of knowledge' (Fricker 2007: 44), thus implying that a competent testifier is one who knows the

proposition that she is testifying to.²⁵ At other points Fricker speaks of ‘the level of credibility’ a hearer attributes to a speaker (Fricker 2007: 19), thus implying that competence is a matter of degree. Both of these interpretations has gained favor with different philosophers. For example, Derek Egan Anderson (2017: 214) favors the former interpretation, while Ole Koksvik (ms.) favors the latter.

Regardless of what Fricker has in mind here, for at least some prominent contextualists (e.g., DeRose 2002; Pynn 2015 & forthcoming), credible testimony must be understood as testimony that is an expression of ‘knowledge’. This is because these contextualists endorse the knowledge norm of assertion:

ASSERTION (A): it is permissible for S to assert that p iff S satisfies ‘knows that p’ in S’s context (DeRose 2002: 182).

When we testify, we assert a proposition to a hearer. So, if a speaker should only assert a proposition when they ‘know’ it, then a speaker should only testify to a proposition when they ‘know’ it. If a speaker should only testify to a proposition when they ‘know’ it, then it is hard to see how the speaker could be testifying credibly if they don’t ‘know’ the proposition that they are testifying to. If one accepts (A), as numerous contextualists do, then one should also accept a knowledge norm of credible testimony:

KNOWLEDGE NORM OF CREDIBLE TESTIMONY (KT): S’s testimony that p is credible iff S ‘knows that p’ in S’s context.

Geoff Pynn is a contextualist (2015) who endorses (A) (forthcoming) and acknowledges that such

²⁵ The interpretation of credible testimony as knowledgeable testimony gains further support from the fact that all Fricker’s central cases of testimonial injustice involve a rejection of speakers’ knowledgeable assertions. See Fricker (2007) chapters 1 and 2 of for a full description of these cases.

an endorsement implies that testimonial injustice should be understood as wrongly not taking a speaker to 'know' a proposition that they testify to because of the operation of a prejudice. He writes: 'violations of the entitlement to assert constitute the fundamental wrong of testimonial injustice' (forthcoming: 17), and that 'an identity-prejudicial credibility deficit perpetrated against a non-knower thus does not degrade its victim in the way characteristic of testimonial injustice' (forthcoming: 18).

A perhaps tempting move for the contextualist to make at this stage would be to insist that credibility comes in degrees and to claim that a testifier is more or less credible depending on how close their epistemic position puts them to the contextually relevant epistemic standard. But this is not a legitimate move for those who endorse (A). What it means for a person to be competent is for that person to have the ability to adequately carry out some tasks. There are some tasks where one's ability to adequately carry them out is a matter of degree, for example, one's ability to play football. It's not that one either has or does not have the ability to play football, rather it's that one can play better or worse. There are other tasks where one either has or does not have the ability to carry it out. If (A) is correct, then the ability to testify is like the latter. A speaker testifies correctly when they do so in accordance with the rules governing the practice of testimony. Testifying in accordance with the rules governing the practice is to testify only to that which you 'know'. Testifying on the basis of anything less is to fail to act in accordance with the rules governing the practice. No matter how close a speaker is to standing in a relation of 'knowledge' to a proposition, they fail to act in accordance with the norms governing the practice of testimony if they testify to a proposition that they do not 'know'. So, if (A) is correct, a testifier is competent iff they 'know' that which they testify to.

Of course, the contextualist could avoid being committed to (KT) by giving up (A) and instead

endorsing some less demanding norm of assertion. However, such a move would come at a significant theoretical cost. Contextualists' endorsement of (A) arms them with the resources they need to respond to one of the most pressing objections to contextualism and provides a key argument in favor of contextualism. Keith DeRose utilises (A) to respond to so called warranted assertability manoeuvres (WAMs) against contextualism. The WAM objection says that while it intuitively seems as though the truth-conditions of knowledge claims vary with context, this is a mistake. Rather, the advocate of WAMs claims, what our intuitions are in fact tracking are the contextually variable warranted assertability conditions of knowledge claims (see, e.g., Brown 2006 for the WAM objection). DeRose (2002) responds with an argument that has (KA) as a premise that at once provides both a response to WAMs and an argument for contextualism. In short, the argument says that if (KA) is true and whether one is warranted in asserting that p varies with the context of utterance, then so too do the conditions for knowledge claims (DeRose 2002: 187).²⁶

So, while it is an option for the contextualist to respond to the dilemma that I have set up by giving up (A), doing so leaves the contextualist facing a further dilemma: if contextualists want to continue endorsing (A), then they must hand over the resources to the prejudiced that enable them to pull off the trick of avoiding the charge of perpetrating testimonial injustice and concede that cases such as **DISMISSED FROM CLASS** do not involve testimonial injustice once certain error possibilities are raised. If the contextualist ditches (A) in order to avoid being committed to (KT), then they must give up both their current response to WAMs and a powerful argument in favor of contextualism.

Some contextualists may be willing to pay the price of rejecting (A). After all, not all contextualists have endorsed (A) as explicitly as DeRose and Pynn have. But there is also a further reason for

²⁶ For a more detailed explanation of this argument see chapter 2, Section 5.

accepting (KT) that is independent of (A), which stems from accepting a particular view in the epistemology of testimony, namely, assurance views of testimony (e.g., Hinchman 2005). According to assurance views of testimony, when S tells H that p, S invites H to trust (or S assures, or guarantees) him that p. If H accepts S's invitation, then she has acquired a reason to believe that p. Proponents of such views rely upon the distinction between telling and mere assertion to explain how it is that H acquires a reason to believe that p. When we tell an audience that p we're not merely uttering that p, rather we're offering them a guarantee or an assurance that p is the case. If the audience accepts this, an interpersonal relationship is established between us and them. This is what gives H a reason to believe that p. H, and H alone, has been assured of the truth of p by S. If the assurance view is correct, then we should accept (KT). Even if you reject that mere assertion is governed by a knowledge norm, we should think that telling is governed by a knowledge norm according to which it is permissible to testify that p iff you know that p. When you're testifying, you're offering a guarantee, or an assurance, or issuing an invitation to trust you that p. You would be criticizable if you offered up such an assurance without knowing that for which you're giving assurance. Such a norm explains the appropriateness of statements like 'You should not have told me that p if you didn't know that p!'

What if (A) and the assurance view of testimony are both rejected? Would that pave the way for contextualists to diagnose **DISMISSED FROM CLASS** as a case of testimonial injustice and avoid the dilemma that I've set up? Let's see. Rejecting (A) and the assurance view allows one to avoid being committed to (KT) and opens up the possibility of contextualists endorsing the view that credibility is a matter of degree. This enables the contextualist to say that there are cases where a hearer perpetrates testimonial injustice against a speaker even though according to contextualism the hearer is correct in denying that the speaker 'knows' the proposition that they're testifying to. For example, a hearer might correctly deny that a speaker 'knows' given that they fall short of the

contextually operative epistemic standard. Nevertheless, the hearer might still perpetrate testimonial injustice by failing to update their beliefs to the degree that they should've because they take the speaker to be slightly less credible than they are owing to the operation of an identity prejudice in the hearer.

However, even though the assumption that credibility is a matter of degree allows contextualists to say that there *can be* testimonial injustice despite knowledge denial in some cases (something the contextualist could not say when assuming that credible testimony is knowledgeable testimony or when assuming an assurance view of testimony), this does not mean that contextualists can always correctly diagnose testimonial injustice where the contextualist verdict is that 'S does not know that p'. This is because there can be cases in which giving the speaker the credibility that they're due requires ascribing knowledge to them. In such cases, treating the speaker as a non-knower because of the operation of a prejudice is to perpetrate testimonial injustice. **DISMISSED FROM CLASS** is an instance of such a case; Rashaan suffers from testimonial injustice because the Assistant Principal denies that he 'knows that Ms. H's decision to send him out was guided by racial prejudice'. What the contextualist needs in order to diagnose **DISMISSED FROM CLASS** as a case of testimonial injustice is to be able to say that Rashaan does 'know that Ms. H's decision to send him out was guided by racial prejudice'. But the contextualist cannot say this, irrespective of whether they assume that credible testimony is always knowledgeable testimony, or that the credibility of testimony is a matter of degree.

4.2. Contextualism is Irrelevant to Epistemology

A second strategy for resisting the dilemma says that once we're clear on what contextualism is a thesis about, we'll see that contextualists aren't committed to saying that **DISMISSED FROM**

CLASS isn't a case of testimonial injustice. Contextualists might point out that their thesis is primarily a linguistic one about the semantics of knowledge ascriptions, and not an epistemic thesis about knowledge itself. Given this, even if contextualism provides the correct semantics for 'knows', it might be irrelevant to the nature and limits of knowledge itself. Testimonial injustice concerns the latter: knowledge itself. It occurs when a hearer prevents a speaker from imparting the property of knowledge. Putting these thoughts together, the contextualist can deny that they face a dilemma. Even though contextualists must say that the Assistant Principal speaks truly in denying the truth of Rashaan's knowledge ascription, they can still say that the Assistant Principal perpetrates testimonial injustice by denying that Rashaan has knowledge itself.

But, as I argued in section 2, testimonial injustice does not only occur when a hearer prevents a speaker from imparting the property of knowledge itself. It can also occur when a hearer wrongfully denies that 'S knows that p'. That is, testimonial injustice can take the form of TI-'Knowledge'. It is this subclass of cases of testimonial injustice that we're interested in. In **DISMISSED FROM CLASS**, Rashaan suffers from TI-'Knowledge' because his Assistant Principal wrongfully denies that he 'knows that Ms. H.'s decision to send him out was guided by racist prejudice'. Contextualism fails to deliver the verdict that Rashaan suffers from TI-'Knowledge' because it says that the Assistant Principal's knowledge denial is true. This is what gives rise to the dilemma. Pointing out the epistemological irrelevance of contextualism does nothing to sidestep this.

What's more, contextualists are unlikely to be tempted to accept that their view is epistemologically irrelevant as doing so would require them to give up much of what's valuable about their view. While it's true that contextualism is a semantic thesis, it's a semantic thesis that contextualists take to do real and substantive epistemological work. One of the biggest selling points of contextualism is its ability to provide a response to skepticism. The kind of skepticism that contextualists are

responding to is skepticism about knowledge itself. If contextualists respond to the dilemma that I've set up by emphasizing the epistemological irrelevance of their view, then they can't maintain that they've provided a response to skepticism. So, unless the contextualist wants to give up their response to the sceptic, then emphasizing the semantic aspect of their view isn't a good response to the dilemma. As Brian Montgomery (2017: 123) has said in his discussion of the epistemological relevance of contextualism: 'it would be disastrous for the epistemic contextualist to sever the link between the two [i.e. the semantics and metaphysics of knowledge]'.

What's more, the only philosophers who emphasize the semantic aspect of contextualism whilst downplaying its significance for epistemology tend to be opponents of contextualism. For example, Ernest Sosa has said contextualism might 'overreach' in its attempt to show that its semantic findings have implications for epistemology (2000: 3). Hilary Kornblith has plainly stated that 'contextualism is irrelevant to epistemology' (2000: 25). Contextualists have been deeply resistant to the idea that contextualism is epistemologically irrelevant. In commenting on Sosa's remarks, DeRose (2017) has polemically said that Sosa plainly gets contextualism 'all wrong'. Whatever their opponents say, it's clear that contextualists wouldn't want to appeal to contextualism's epistemological irrelevance in response to the dilemma I've set up.²⁷

4.3. High Standards Don't Always Prevail

A third way for the contextualist to try to resist the dilemma is to claim that it is a mistake to think that contextualism is committed to the view that when conversational participants are pushing the epistemic standards in different directions, that the higher standards always prevail. When the

²⁷ Some are explicitly resistant to making such an appeal. In addition to DeRose (2017) see also Blome-Tillmann (2007) for an additional response to Sosa's (2000) claim that contextualism is epistemologically irrelevant.

Assistant Principal mentions the possibility that Rashaan has misinterpreted Ms. H's actions, he is met with an 'Of course, it's a possibility! But, I've seen white kids not get sent out for doing the same thing!' from Rashaan. Rashaan protests against the Assistant Principal's attempt to put high standards in place by appealing to a lower standard where Rashaan's observational evidence is sufficient for his knowledge claim to be true. Whose standards prevail? If it can be shown that the higher standards of the Assistant Principal don't win out, then the claim that 'Rashaan knows that Ms. H's decision to send him out was guided by racial prejudice' will not be false, and the contextualist may have a way out of the dilemma. Different contextualists have different answers to the question of how epistemic standards are determined. Here, I will consider DeRose's gap view, discussed in chapter 3, and Blome-Tillmann's presuppositional epistemic contextualism, discussed in chapter 4, and show that both result in the dilemma. We will see that the same features of the views that led to the problems identified in chapters 3 and 4 also mean that they cannot account for our intuitions about cases of testimonial injustice.

Firstly, Keith DeRose (and Stewart Cohen).²⁸ As we know, DeRose subscribes to what he calls the gap view of how epistemic standards are determined in situations where the personally indicated standards of conversational participants diverge. In DeRose's own words, here is a quick reminder of the gap view:

Applied to cases involving 'know(s)', then, in cases of small divergence in personally indicated epistemic standards, 'S knows that p' is true (and 'S doesn't know that p' is false) where S meets the personally indicated standards of both speakers; 'S doesn't know that p' is true (and 'S knows that p' is false) where S fails to meet either set of standards; and where

²⁸ DeRose points out that other contextualists share this view. He writes: 'I should report that when I asked Stewart Cohen about the situations under consideration, he too was inclined to say that neither the skeptic nor her opponent was speaking the truth, so he is probably inclined to something like the gap view...' (2009: 148).

S meets one set of standards but fails to meet the other, both ‘S knows that p’ and ‘S doesn’t know that p’ go truth-value-less (DeRose 2009: 144-145).

According to the gap view, when the personally indicated standards of conversational participants diverge, each of the participants’ standards register in determining the operative standard. In order for the knowledge ascription to be true the standards of both conversational participants must be met. Both standards must fail to be met if the knowledge ascription is to be false. If one standard is met while the other standard is not met, then the knowledge ascription goes truth-valueless.

The gap view might look promising. The Assistant Principal’s standards don’t automatically prevail and Rashaan’s personally indicated standards contribute towards determining the operative epistemic standard. Unfortunately, the gap view doesn’t vindicate our intuitive verdict on **DISMISSED FROM CLASS**. For **DISMISSED FROM CLASS** to instantiate testimonial injustice it must be true that: (1) Rashaan satisfies ‘knows that Ms. H’s actions were guided by racial prejudice’; (2) Rashaan testifies that Ms. H’s actions were guided by racial prejudice, and; (3) the Assistant Principal is motivated by identity prejudice in rejecting Rashaan’s testimony. While the truth of (2) and (3) are secure, the truth of (1) isn’t. For (1) to be true on the gap view Rashaan must meet both his own epistemic standards, and the much higher standards of the Assistant Principal. While Rashaan meets his own standards, he fails to meet those of the Assistant Principal. The verdict of the gap view is that (1) is truth-valueless. The truth of (1) is thus not secured and **DISMISSED FROM CLASS** does not instantiate testimonial injustice.

The feature of the gap view that means it cannot vindicate our intuitions about **DISMISSED FROM CLASS** is the same as the feature that makes it vulnerable to trolling (chapter 3), namely, that the personally indicated standards of every conversational participant registers, including that of the Assistant Principal. So, there are similarities between the problem posed by cases like **DISMISSED**

FROM CLASS and cases of trolling. The two classes of problem cases are structurally similar in that they are conversations that involve one conversational participant having different epistemic standards from another conversational participant which leads to the undermining of an intuitively true knowledge claim. The most important difference is that testimonial injustice cases, like **DISMISSED FROM CLASS**, necessarily involve prejudice, which is not true of trolling cases. It's important to recognize this as a distinct problem because of the importance of being able to diagnose testimonial injustice. Note also that, unlike trolling cases, **DISMISSED FROM CLASS** is not a case where someone just walks into the conversation to derail the goals of their interlocutor.

Let us turn now to Michael Blome-Tillmann's presuppositional epistemic contextualism (PEC) (2009, 2014), which is perhaps the most sophisticated contextualist attempt to handle skeptical worries. Recall from chapter 4 that Blome-Tillmann denies that raising an alternative to salience automatically makes it relevant. Rather, on Blome-Tillmann's account, whether a raised error possibility is relevant depends on what conversational participants pragmatically presuppose:

Pragmatic Presupposition (PP): S pragmatically presupposes p in C iff S is disposed to behave, in her use of language, as if she believed p to be common ground in C (Blome-Tillmann 2014: 26, amended from Stalnaker 1974).

Where common ground is to be understood as follows:

Common Ground (CG): It is common ground that p in a group G iff all members of G accept (for the purposes of the conversation) that p, and all believe that all accept that p, and all believe that all believe that all accept that p, and so on (Stalnaker 2002: 704).

According to PEC, if an error possibility is mentioned, but is not compatible with conversational participants' pragmatic presuppositions, then it is not added to the set of relevant alternatives. So,

according to PEC, 'S knows that p' is true in context C if S's evidence eliminates all of the not-p possibilities that are compatible with what is pragmatically presupposed in C (Blome-Tillmann 2014: 340).

Applied to **DISMISSED FROM CLASS**, PEC says that even though the Assistant Principal mentions the possibility that Rashaan has misinterpreted Ms. H's actions, a skeptical context is not automatically put in place. Clearly Rashaan does not pragmatically presuppose that it is possible that he is misinterpreting the actions of others that are directed at him such that he sees racism where there is none. If the Assistant Principal recognizes this, he may cease to insist on it, and so cease to be disposed to behave as if it is part of the common ground and as if it needs to be ruled out. If the Assistant Principal did so cease to insist, then we would have a contextualist resolution to the dilemma. However, problems arise for PEC when a conversational participant insists on a skeptical possibility. In order for it to continue to be true that a subject 'knows' after being confronted with a skeptical error possibility, the negation of the skeptical error possibility (not-SEP) must be mutually pragmatically presupposed by all participants in a conversation (2014: 47). If the skeptic refuses to pragmatically presuppose not-SEP, and instead continues to insist on a skeptical error possibility then, Blome-Tillmann says, we end up in a 'defective context'; a context in which it is unclear whether 'knows' is satisfied (2014: 43 - 45). Clearly this does not generate the desired verdict in **DISMISSED FROM CLASS**. So long as the Assistant Principal continues to insist on the skeptical error possibility, a vindication of the intuition that 'Rashaan knows that Ms. H's actions in sending him out of class were guided by racial prejudice' is not forthcoming, and the dilemma remains. This parallels how PEC deals, or fails to deal with, the **TEENAGER** case discussed the previous chapter. It deals with the case only when the skeptical speaker is co-operative, and fails to deal with it otherwise.

Recall from chapter 4 that Blome-Tillmann discusses a strategy that might be employed to ensure that one is in a context where not-SEP is mutually pragmatically presupposed. He says that if we find ourselves in a context where someone is insisting on a skeptical error possibility, we might appeal to a ‘conversational authority’ to settle the issue of whether the high standards prevail. In such a situation, conversational authorities might put ‘their foot down’ and declare that the skeptical error possibility is to be ignored (2014: 47 - 48). However, as we saw in the previous chapter, without a well worked out account of who counts as a conversational authority, such an appeal doesn’t look to be particularly helpful. This is especially true of **DISMISSED FROM CLASS**; when it comes to conversations between an Assistant Principal and a student, surely the Assistant Principal’s word is authoritative. This parallels how PEC struggles to deal with authority problem cases, such as **SEXUAL HARASSMENT***, that were discussed in the previous chapter.

I worked through three ways that the notion of conversational authority might be developed in chapter 4 and showed that none of them are satisfactory. I’ll work through just one of those options again here in order to see whether it fares better in the context of trying to find a contextualist resolution to the dilemma we’re presently considering. Remember that one way to flesh out the notion of a conversational authority that could be helpful to the contextualist is to regard a conversational authority as an epistemic authority, in the sense developed by Linda Zagzebski. On Zagzebski’s account a subject is an epistemic authority for you when they are better than you at getting to the truth of some proposition (Zagzebski 2013), perhaps because they have many more reasons pertaining to the issue than you, and because they are more skilled at evaluating those reasons (Dormandy 2018: 773 - 774). If a conversational authority is understood in this way, then Rashaan might be thought to be a conversational authority when it comes to the subject of racism. This is because Rashaan frequently experiences patterns of racial discrimination. Rashaan may thus put his foot down and authoritatively declare that the skeptical error possibility raised by the

Assistant Principal is to be ignored. It would then be true that Rashaan ‘knows that Ms. H’s actions were guided by racism’, and so the Assistant Principal’s knowledge denial would constitute testimonial injustice.

This attempt to flesh out the notion of a conversational authority as an epistemic authority is equally unhelpful in diagnosing cases of testimonial injustice as it was in dealing with authority problem cases. Firstly, it’s not clear that Rashaan is an epistemic authority in the relevant sense. Remember that Zagzebski’s account of epistemic authority is an account of subjective epistemic authority. It is an account of when a subject is an epistemic authority *for you*, not an account of absolute epistemic authority or epistemic authority simpliciter. For Zagzebski, for someone to be an epistemic authority for me requires that I trust that she is more epistemically conscientious than I am myself (2013: 299). For Rashaan to be an epistemic authority for the Assistant Principal requires that the Assistant Principal trusts that Rashaan is more epistemically conscientious than he himself is. But, this trust looks to be absent in **DISMISSED FROM CLASS**. The Assistant Principal has very skeptical inclinations and he possesses an undefeated defeater for Rashaan’s claim to ‘know that Ms. H’s decision to send Rashaan out of her class was guided by racial prejudice’. It’s hard to see how the Assistant Principal can be said to trust that Rashaan is more epistemically conscientious than he himself is when he is a person of skeptical nature who has an undefeated defeater for his interlocutor’s knowledge claim.

Secondly, as discussed in chapter 4, even if we suppose that Rashaan is an epistemic authority for the Assistant Principal, it’s unclear why we should take someone’s being an epistemic authority in Zagzebski’s sense to also be an authority on what the contextually operative epistemic standards should be. As I said in chapter 4, it simply doesn’t follow from the fact that someone both has more reasons than me bearing on the truth of p, and is good at evaluating those reasons, that they have

authority over what the standard of justification should be. The purported epistemic authority's reasons might have a bearing on my beliefs and where I stand in relation to the contextually operative standard of justification, but it is not obvious why they should determine the standard of justification itself.

5. Summary and Preview

In this chapter we've seen that contextualism's failure to vindicate our intuitions about some cases of testimonial injustice – in particular, some cases of 'TI-Knowledge' – presents contextualists with a dilemma. Contextualists must either reject conversational contextualism in its current forms in order to vindicate intuitions about cases of testimonial injustice, or, accept contextualism, hand over the resources that allow the prejudiced to being charged with perpetrating testimonial injustice, and concede that such cases are not in fact examples of testimonial injustice once epistemic standards are raised.

We've come to the end of the first payoff from theorizing conversational contextualism non-ideally: paying attention to the ways in which power effects who can pull off standards raising moves has revealed unrecognized problems with extant versions of conversational contextualism. In chapter 3 we saw that DeRose's gap view is vulnerable to trolling, and that the central reason for this is that DeRose insists that every conversational participant's personally indicated standards must register, and that includes the trolls. In chapter 4 we saw that Blome-Tillmann's presuppositional epistemic contextualism is vulnerable to counterexamples in the form of authority problem cases and adaptive presupposition cases, and that the central reason for this is that the success of Blome-Tillmann's view requires co-operation, but co-operation is not always forthcoming, and that even when it is, the terms of co-operation can be unfair. In this chapter we've seen that all of the extant versions of

conversational contextualism that we've considered cannot vindicate intuitions about paradigm cases of testimonial injustice. There's a pattern: conversational contextualism faces a series of problems because of failures to place adequate constraints on how conversational dynamics affect the truth-conditions of knowledge ascriptions. This failure is at least in part a result of the contextualist having overlooked certain features of the real world: that power asymmetries structure conversations, that co-operation amongst conversational participants can be absent, and that co-operation can be present yet take place on unfair terms. Thus, a view that is supposed to be purely descriptive, and so non-ideal, turns out to rest on some idealizing assumptions. What's the upshot? In order to develop an acceptable contextualism, one that avoids the problems that have been adduced over the course of the last three chapters, we must either find a way to adequately constrain the effects of conversational dynamics on the truth-conditions of knowledge ascriptions, or we must divorce the truth-conditions of knowledge ascriptions from conversational dynamics altogether and instead tie them to something extra conversational. In chapters 7 and 8 I'll run a version of each of these strategies. But, before we get to that, there is the second payoff to get to. Whereas in this chapter I argued that existing forms of conversational contextualism cannot diagnose certain cases of epistemic injustice, in the next chapter I'll argue that theories of conversational contextualism helps us to identify an unrecognized form of epistemic injustice that I'll call irrelevance injustice.

Chapter 6: The Politics of Relevant Alternatives

1. Introduction

By the end of chapter 5 things aren't looking good for conversational contextualism: paying attention paid to how power influences who can pull off relevant alternative expanding (or, standards raising) moves revealed unrecognized problems with existing forms of conversational contextualism. This made up the first payoff of theorising contextualism non-ideally. At the end of chapter 5 I identified the common problem underlying the more specific problems facing particular versions of conversational contextualism, mentioned the options for going forward, and promised to explore those options in chapters 7 and 8. This will make up the third payoff of theorising contextualism non-ideally. But before we get to that, there is the second payoff to deliver on. In this chapter I'll do just that by showing that theorising contextualism non-ideally helps us to identify an unrecognized form of epistemic injustice that I'll call irrelevance injustice. I will develop the theory of irrelevance injustice on the relevant alternatives model of contextualism (e.g., Lewis 1997, Blome-Tillmann 2009).²⁹

We start out this chapter with the motivating thought of the entire thesis: contextualists have neglected the fact that the conversational contexts are embedded within broader social contexts, and facts about these broader contexts leak into conversational contexts to affect which conversational participants can pull off the conversational moves that expand the set of relevant alternatives.³⁰ This chapter further investigates the relationship between the power relations that

²⁹ I might just as well have developed the theory of irrelevance injustice on the standards model of contextualism. See chapter 1 for an explanation of the relation, differences, and ways of translating between the two. I've opted to theorise about irrelevance injustice on the relevant alternatives model because most of the discussion of context shifting mechanism, which is of central importance for me, is cast in these terms.

³⁰ Johnathan Jenkins Ichikawa (2020) is one very recent exception to this neglect of the relationship between contextualism and power. I discuss the connections between Ichikawa's view and my own at length later in the chapter.

pervade our social context and the ability to successfully raise alternatives in order to reveal an unrecognised form of epistemic injustice: irrelevance injustice. Irrelevance injustice occurs either when a speaker raises an alternative that is not taken seriously when it should be, or when a speaker raises an alternative that is taken seriously when it should not be. Irrelevance injustice is a matter of which alternatives are perceived to be relevant and of how these perceptions impact patterns of knowledge ascriptions in ways that are unfair. Asymmetries in whose alternatives are taken seriously affect how much epistemic labour one must undertake in ruling out alternatives prior to being ascribed knowledge. Because knowledge ascriptions have socially valuable functions – as described in chapter 2, they’ve been associated with identifying reliable informants (Craig 1990), bringing inquiry to an end (Kelp 2011), and tracking the norms of assertion (Williamson 2000) and action (Hawthorne & Stanley 2008) – asymmetries in whose alternatives are taken seriously mean asymmetries in who gets to do socially valuable things with knowledge ascriptions. Thus, in addition to illuminating the underexplored relationship between power, raising alternatives, and the perceived relevance of alternatives, this chapter further illuminates the underexplored relationship between epistemic injustice and knowledge ascriptions.

A quick clarification of my aim before we get going. As we will see, identifying the conditions under which irrelevance injustice takes hold has implications for relevant alternatives contextualism in that it sheds light on when an alternative is relevant. However, my primary aim in this chapter is not to explore the consequences of irrelevance injustice for contextualism, but rather to show that

Prior to Ichikawa, Evelyn Brister (2009, 2017) comes closest to having explored this relationship. Drawing on insights from feminist epistemology, Brister argues that philosophical skepticism unfairly concentrates epistemic power in the hands of skeptics, and that this fact of unfairness is a sufficient reason for resisting skeptical moves. Contextualists thus have a way to make it difficult for skeptics to initiate skeptical contexts. While Brister highlights that skeptics unfairly possess epistemic power, she does not explore the possibility that the power to pull off skeptical moves co-varies with the power that attends to one’s social position, and that this may be a site of injustice. Louise Antony (1995) also notes that features of the social milieu in which an alternative is raised influences which alternatives are taken seriously. But Antony’s discussion is not related to epistemic contextualism.

contextualism has the theoretical resources to illuminate an unrecognized form of epistemic injustice.

I proceed as follows: Section 2 highlights the important features of relevant alternatives contextualism given my aims in this chapter. Section 3 examines paradigm cases, and develops a theory, of irrelevance injustice. Section 4 contrasts irrelevance injustice with other forms of epistemic injustice, namely, Fricker's (2007) model of testimonial injustice and Ichikawa's (2020) recently proposed concept of contextual injustice. Section 5 concludes.

2. Relevant Alternatives Contextualism

Consider again the pair of cases that we saw in chapters 1 and 3. I've changed the labels from **LOW** and **HIGH** to **NARROW** and **WIDE** to better reflect the nature of the shift in context emphasized by relevant alternatives contextualists:

NARROW: Hannah and Sarah are discussing whether the bank is open on Saturday. Hannah recalls, "I've been in there on a Saturday before". She says, "I know that the bank is open on Saturday". Sarah agrees.

WIDE: Hannah and Sarah are discussing whether the bank is open on Saturday. Hannah says, "I've been in there on a Saturday before". But Sarah replies "Banks can change their opening hours. You don't know that the bank is open on Saturday". Hannah agrees. (variation on McKenna 2017: 1, variation on DeRose 1992).

Again, we're assuming that Hannah is in the same epistemic position in both cases and that the bank is open. Intuitively, both Hannah's knowledge ascription and Sarah's knowledge denial are true. The set of alternatives to the proposition that the bank is open on Saturdays that's relevant in

NARROW is narrower – it doesn't include the possibility that the bank has changed its opening hours – than the set of alternatives that are relevant in WIDE – which does include the possibility that the bank has changed its opening hours. Hannah's evidence rules out the set of relevant alternatives in NARROW, but not those that are relevant in WIDE. Thus, it's true to say that Hannah 'knows' in NARROW but not in WIDE.

Here are the important points about relevant alternatives contextualism to keep in mind going forward in this chapter: Firstly, contextualism is a view about knowledge ascribing sentences – sentences of the form 'S knows that p' – and knowledge denying sentences – sentences of the form 'S does not know that p', not about the metaphysical property of knowledge itself. The injustice that I'll be describing involves knowledge ascriptions and denials, rather than knowledge itself. Secondly, we can make mistakes in our judgements about the truth or falsity of knowledge ascribing and knowledge denying sentences. 'S knows that p' might be true, yet someone might mistakenly judge it to be false. 'S knows that p' might be false, yet someone might mistakenly judge it to be true. As we'll see, irrelevance injustice involves these kinds of mistaken judgements about the truth and falsity of knowledge ascriptions and denials. Thirdly, one cause of mistaken judgements about the truth-values of knowledge claims can be mistaken judgements about whether alternatives are relevant or irrelevant. When a speaker raises an alternative by saying 'what about this not-p alternative?', it might be judged to be relevant when it is irrelevant, leading to mistaken knowledge denials, or the alternative might be judged to be irrelevant when it is relevant, leading to mistaken knowledge ascriptions. As we'll see, irrelevance injustice involves these kinds of mistaken judgements about the relevance and irrelevance of raised alternatives, which in turn lead to mistaken judgements about the truth-values of knowledge ascriptions and denials.

What makes an alternative relevant? This is a difficult and much debated question (see, e.g., Austin 1946, Dretske 1970, Stine 1976, Goldman 1976, Lewis 1996, Blome-Tillmann 2009 for discussion). David Lewis's attempt to specify a list of rules that tell us when an alternative is relevant is the most comprehensive attempt to answer this question. Lewis's rules include the *rule of actuality* which says that if some not-p alternative is actual, then it is a relevant alternative (1996: 554 - 5), the *rule of belief* which says that alternatives that the subject in the knowledge claim believes or should believe are relevant (1996: 555), and the *rule of resemblance* which says that alternatives that saliently resemble other relevant alternatives are themselves relevant (1996: 559).³¹ What's important for our purposes is that speakers can change which alternatives are relevant by making certain conversational moves. Which conversational moves affect which alternatives are relevant is a matter of dispute. Lewis (1996: 559) proposed the *rule of attention* which says an alternative becomes relevant when a conversational participant brings it to the attention of other conversational participants. For example, when Laura raises the possibility of the bank changing its opening hours it becomes relevant. The truth-conditions for the ascription change such that the possibility of changed hours must be ruled out if the knowledge claim is to be true. As explained in chapter 4, this rule is now widely rejected because it makes 'knows' too difficult to satisfy. It allows that I can always undermine the truth of knowledge ascriptions by raising a far-out sceptical hypothesis (see, e.g., Williams 2001: 15, Blome-Tillmann 2009: 246, Ichikawa 2017: 25 for discussion of this criticism). In section 3.b. I'll explain that another reason to reject the rule of attention is that it leads to misdiagnoses of cases of irrelevance injustice. As was also discussed in chapter 4, Blome-Tillmann (2009: 249 - 256) proposes a different view on which in order for an alternative to be relevant not only must attention be drawn to it, but it must also be taken seriously by the participants

³¹ There are yet further rules from Lewis that we might want to endorse. For example, he considers the possibility of endorsing a *rule of high stakes*, which says that a high number of alternatives are relevant when error would be especially costly (Lewis 1999: 556, fn. 12).

in the conversation. Ichikawa (2015) has argued that this rule is subject to counterexamples. In the next section I'll make some of my own suggestions for what makes an alternative relevant. In particular, I'll suggest that if an alternative is raised in an exercise of intellectual virtue, then it is relevant. This suggestion will be developed at length in chapter 7. I'll also suggest that if an alternative is taken seriously by someone who is competent in a domain, then it is relevant. At this stage in the thesis, these suggestions are motivated and supported by their ability to help us to identify cases of irrelevance injustice and develop a theory of irrelevance injustice.

Let me reiterate the important points to take from this section into the rest of the chapter: relevant alternatives contextualism is the view that the truth of 'S knows that p' requires that S rule out different alternatives in different contexts. Which alternatives are relevant can be a matter of conversational participants making certain conversational moves. People can make mistakes about the relevance and irrelevance of alternatives. These mistakes can cause further mistakes about the truth or falsity of knowledge ascriptions and denials. The project of the rest of the paper is to show that these mistakes can be a site of injustice.

3. Irrelevance Injustice

Running parallel to conversations among contextualists about how speakers can affect which alternatives are relevant have been conversations about how speakers' social identities and the power that attends to them affect which conversational moves a speaker can pull off. The literature on epistemic injustice, and testimonial injustice in particular, has been one important site of discussion. Canonically, Miranda Fricker (2007) argues that identity power operates in testimonial exchanges to produce testimonial injustice. Fricker defines social power as a capacity to control others' actions (2007: 13). Identity power is social power that depends on shared imaginative

conceptions of a given identity (2007: 14 – 15). In cases of testimonial injustice, a hearer has an identity prejudice which operates as identity power to prevent a speaker from conveying knowledge (2007: 28). Amongst the many others to discuss this theme include Karen Jones (2002), Kristie Dotson (2011), Gaile Pohlhaus Jr. (2012), Jose Medina (2013), and Derek Egan Anderson (2017). Another important site of discussion is the silencing literature. The core idea is that speakers can be silenced by being made unable to pull off certain speech acts. For example, Rae Langton (1993) famously argues that women are silenced in their attempts to refuse sex as they are unable to secure uptake for their locutions from men because of the effects of pornography on men's ability to recognise women's intention to refuse sex. Amongst the many others to discuss this theme include Rae Langton & Jennifer Hornsby (1998), Rae Langton & Caroline West (1999), Ishani Maitra (2009), Mary Kate McGowan (2009), Alex Davies (2016), and Samia Hesni (2018).

Largely missing from these discussions (Ichikawa (2020) provides the one very recent exception which I'll discuss at length in section 4) is an investigation into how the two conversations interact. In this section I investigate how social identity and power affect whether speakers can successfully raise alternatives as relevant, how this affects the practice of ascribing and denying knowledge, and whether injustice lies therein. A distinctive form of epistemic injustice is revealed: irrelevance injustice.

I'll proceed by examining three cases of irrelevance injustice before arriving at a definition. Each case demonstrates a different form that irrelevance injustice can take. Along the way I highlight points of difference from already recognised forms of epistemic injustice and some implications for how to understand epistemic injustice. Here's the first case:

BOARDROOM: The employees of Punch are meeting to discuss the company's recent decline in profits. Mr X asserts, 'I know that spending inefficiencies are the sole explanation

of the company's decline in profits.' Miss Triggs, a competent business analyst and the only woman present, points out that the economy has recently undergone structural changes, and that if Mr X can't rule that out as an explanation, then he does not know that profits have declined solely because of spending inefficiencies. The structural changes explanation is relevant and is not ruled out by Mr X's evidence. Yet, because they harbour a prejudice that women are not well placed to make claims about business and economics, the other meeting participants dismiss Miss Triggs' alternative as irrelevant and continue to ascribe knowledge of the spending inefficiency explanation to Mr X.³²

BOARDROOM is a case of irrelevance injustice. Miss Triggs raises a relevant alternative that is not taken seriously by her more powerful audience because of a prejudice they harbour against women that results in direct epistemic harm to her and in distortions in the pattern of knowledge ascriptions. The alternative is relevant because Miss Triggs raises it in an exercise of her competence as a business analyst. This is in line with Georgi Gardiner's (2020: 12) claim that it is dogmatically stubborn to ignore error possibilities that other competent people take seriously. A further reason why Miss Triggs's alternative is relevant is because she raises it in an exercise of the intellectual virtue of open-mindedness. As the disposition to consider a variety of viewpoints in order to further one's understanding or to get at the truth, open-mindedness is the intellectual virtue that's to be most closely associated with raising relevant alternatives (see Riggs 2019 for an overview of open-mindedness). Miss Triggs plausibly displays open-mindedness when she raises the alternative that Punch's decline in profits is at least partly explained by structural changes in the economy. I'll say more about why these factors make an alternative relevant in chapter 7 where I

³² **BOARDROOM** is inspired by Riana Duncan's 1988 Punch magazine cartoon depicting a sexist boardroom. The cartoon shows five men and one woman sitting down to a meeting in which the woman, Miss Triggs, makes a suggestion that's quickly dismissed. Underneath the cartoon the caption reads 'That's an excellent suggestion, Miss Triggs. Perhaps one of the men here would like to make it'.

propose a version of contextualism according to which alternatives that are raised in acts of intellectual virtue are relevant. These facts mean that the alternative is relevant even if we suppose that structural changes in the economy are no part of the explanation for Punch's recent misfortunes, and so are not ruled in by being actual.

Miss Triggs is not taken seriously by her more powerful audience. The relevant form of power here is *epistemic power*. Kristie Dotson (2018: 130) defines epistemic power as 'differing ranges of privilege and underprivilege that's co-extensive with one's epistemic status' (2018: 130). She defines *epistemic status* as 'positive or negative assessments of one's epistemic position' (2018: 131). Epistemic status is domain specific and can be legitimate, such as when positive epistemic status is conferred on a climate scientist with respect to climate change, or illegitimate, such as when positive epistemic status is conferred on a man with respect to business simply because he's a man. The privileges conferred by having high epistemic status include the ability to have one's claims taken seriously and being able to dismiss the claims of others. In Dotson's terms, having epistemic power enables one to seize the *epistemological high ground* which is a 'contrastive privilege that generates defense of one's claims and the ability and/or authority to challenge competing claims' (2018: 139).

Applied to BOARDROOM, the men have epistemic power, in that they enjoy the positive epistemic status of being judged competent in the world of business because of the idea that men are in the know in the world of business. This gives them the privilege of being able to dismiss Miss Triggs's (relevant!) alternative and subsequently continue being (falsely!) ascribed knowledge.

A helpful way to see that it's not the mere failure to take an alternative seriously that makes for irrelevance injustice, but rather failure to take an alternative seriously by those who possess power in the context, is by considering what happens when those who *lack* power fail to take an alternative seriously. Suppose that it's Mr X who raises an alternative, and that it's not taken seriously by Miss

Triggs. It's hard to imagine, given her social identity and attendant lack of social and epistemic power in the context, that Miss Triggs not taking Mr X seriously would prevent his alternative from being judged to be relevant and in need of ruling out before knowledge is ascribed.

Now to the harms and consequences. In being taken to lack the epistemic competence required to contribute to the investigation at hand, Miss Triggs suffers an affront to her epistemic agency that is the hallmark of epistemic injustice. This affront plausibly has further harmful consequences. Epistemically, Miss Triggs may suffer from diminished intellectual courage – the intellectual virtue of not ceasing to contribute to an epistemic activity too easily – due to having her confidence shaken by being so swiftly dismissed. Practically, Miss Triggs might miss out on professional benefits that come with being seen as a meaningful contributor to a workplace meeting.

The harms to Miss Triggs highlight one of several important differences between irrelevance injustice and testimonial injustice; the injustice in which a speaker is not given the credibility they are due when testifying because of a hearer's identity prejudice. According to Fricker, the victim of testimonial injustice is harmed in their capacity 'as a giver of knowledge' (2007: 45). But whether Miss Triggs 'knows' is never at issue. In raising an alternative, she is not attempting to give knowledge, rather she's attempting to take part in an inquiry into a particular question. Irrelevance injustice thus supports Mikkil Gerken's (2019: 2) proposed broadening of epistemic injustice from involving a harm done to someone in their capacity as a knower, as Fricker proposes (2007: 1), to harm done to someone in their capacity as an epistemic subject.

While the harms suffered by Miss Triggs are significant, the downstream effects are perhaps even more significant. Dotson (2018) points out that epistemic power can be accumulated, which can in turn serve to uphold and exacerbate unjust power relations that set up the conditions for the perpetration of further injustice. **BOARDROOM** provides a stark illustration of this process of the

accumulation of epistemic power. In irrelevance injustice alternatives that are in fact relevant and alternatives that are taken to be relevant come apart in a way that leads to distortions in the patterns of knowledge ascriptions that facilitate the accumulation of epistemic power and the stabilization of injustice. Because Miss Triggs's alternative is relevant, Mr X's knowledge ascription is literally false. But because those with epistemic power don't take Miss Triggs seriously, they mistakenly judge that Mr X's knowledge ascription is true. This is the distortion: Mr X is ascribed knowledge when he should not be, and the merely apparent truth of Mr X's knowledge ascription is sustained by the judgements of those with epistemic power.

The significance of this distortion is tied up with the socially valuable functions of knowledge ascriptions. As we know from chapter 2, it has been suggested that knowledge ascriptions function to identify reliable informants (Craig 1990), signal the appropriate end of inquiry (Kelp 2011), and track the epistemic norms governing assertion (Williamson 2000), belief, and practical reasoning (Hawthorne & Stanley 2008). In being ascribed knowledge when he should not be, Mr X accumulates greater epistemic power by maintaining statuses and permissions that he is not entitled to: he's seen as a reliable informant, as someone who has reached the end of inquiry, and as someone whose epistemic position permits him to assert and act on the proposition. In these ways, Mr X benefits from the injustice Miss Triggs faces. These benefits serve to stabilize unjust power relations that set up the conditions for further injustice. Bearing the positive epistemic status of a reliable informant and as a place where inquiry can legitimately stop means that the other conversational participants will be led to misplace their trust in Mr X. Such misplaced trust causes the other conversational participants to be misled both about who 'knows' and about what is 'known'. This consolidates power in the hands of Mr X, as in the eyes of his audience he is a reliable informant who is well placed to dismiss challenges to his knowledge claim. When the further relevant conditions obtain, such dismissals constitute further injustices.

How much power must an audience have if their reception of an alternative is to have the harmful consequences characteristic of irrelevance injustice? It's hard to provide a precise answer, but we can at least say that the harms of irrelevance injustice will be most stark when there are serious power differences between the audience of the alternative and the victim of irrelevance injustice. Things will be more complicated when there is a less dramatic power differential. Some but not all of the harms might take hold, or the harms may be less severe.

In **BOARDROOM**, the men do not take Miss Triggs seriously because of a prejudice that they hold against women. A variation on **BOARDROOM** shows that prejudice is not a necessary condition of irrelevance injustice. Imagine a version of the case in which everything is the same apart from the reason why the men do not take Miss Triggs seriously. Imagine that instead of not taking Miss Triggs seriously because of prejudice that women are not well placed to make claims about business and economics, they instead fail to take her seriously because they feel threatened by her intelligence. In this version of the case the men use their epistemic power to work together to make it appear as if Miss Triggs's alternative is irrelevant and thereby undermine her ability to take part in the discussion on an equal footing. The reason that this is an injustice despite the absence of prejudice is that the judgement of irrelevance and the harmful consequences that flow from it conform to a general pattern of oppression that unfairly harms some and privileges others in the same ways as those judgements that are caused by prejudice. Judgements of relevance that conform to the general pattern of norms that work to harm the oppressed by making it difficult for them to be ascribed knowledge and privilege the dominant by making it easier for them to be ascribed knowledge, are of concern even when conformity does not arise from prejudice. Following

the Fricker paradigm in which negative identity prejudice is a necessary condition for epistemic injustice will cause us to miss out on cases like this.³³

Consider a second case, a variation on **DISMISSED FROM CLASS** that was our principal no testimonial injustice case in Chapter 5

CLASSROOM: During class Rashaan, a Black public high school student, gets up out of his seat numerous times to get a tissue from the tissue box. His teacher, Ms. H., interprets his behaviour as disruptive and sends him to the Principal's office. Ms H's decision to send Rashaan out was guided by racist stereotypes of black students as troublemakers.³⁴ Rashaan has witnessed non-Black students behave just as he did in Ms. H's class and yet not get sent out. From this he justifiably infers that Ms. H's actions were guided by racism. When Rashaan reaches the Principal's office, he finds the Principal in conversation with the Assistant Principal. Both are white men. They're discussing the Principal's recent experience of taking his first anti-racism course. The Principal stops his conversation, invites Rashaan in, and they begin to chat:

P: Rashaan, why have you been sent to my office?

R: I was just collecting tissues to blow my nose. I know that this is racism!

³³ See Derek Anderson (2017: 211) for further discussion on this point.

³⁴ Given the findings of the studies that were described in footnote 24, it is reasonable to suppose that it is true that Ms. H's action is guided by a stereotype of black students as troublemakers. Okonofua and Eberhardt (2015) show that teachers are likely to interpret students behaviour differently based on race: teachers are more likely to perceive minor infractions by Black students as being more severe than those of white students; report being more irritated by infractions of Black students rather than white students; think it appropriate to punish Black students more harshly than white students for equivalent infractions, and; are more likely to view Black students, rather than white students, as troublemakers. The interaction between Rashaan and Ms. H. is based on one of the cases that Okonofua and Eberhardt give to participants in their study that elicited harsher responses when the behaviour was said to be carried out by a black student as compared to a white student. They based the case on real office referral reports collected from California public middle schools (2015: 618).

P: Rashaan, it's possible you're mistaken. Can you rule out the possibility that you've misinterpreted Ms. H's actions? It's possible that you're systematically misinterpreting the actions of others when directed at you!

AP: What the Principal says is right!

R: That's irrelevant! I've seen white kids not get sent out for doing the same thing!

P: If you can't rule out the possibility that you've misinterpreted Ms. H's actions, then you don't know that Ms. H's actions were guided by racism.

The Principal, confidence bolstered by his recent training, and Assistant Principal, convinced of the effectiveness of the training, judge the Principal to be capable of detecting racism. With Rashaan unable to eliminate the possibility that he's misinterpreted Ms. H's actions - a possibility that he cannot rule out with his own observational evidence alone - the Principal denies that Rashaan knows that Ms. H's actions were guided by racist prejudice.

CLASSROOM is also a case of irrelevance injustice, but in a different way from BOARDROOM. Whereas BOARDROOM involves the dismissal of a relevant alternative because of a prejudice against the speaker, in CLASSROOM the Principal raises an irrelevant alternative that is taken seriously, by those with power in the context, because of a prejudice they harbour in the Principal's favour.³⁵ The Principal's alternative is irrelevant, and Rashaan is competent to judge it to be. Above I suggested that an alternative is relevant if it is raised by someone who is competent in the relevant domain. We can take this idea further by saying that we can also figure out if an alternative is

³⁵ The injustice would be compounded if the Principal's motivation for raising the skeptical alternative was racist prejudice. For example, a prejudicial belief that black people often see racism where there is none. But this is not essential to the perpetration of irrelevance injustice. What is essential to this form of irrelevance injustice is that there's a prejudice in favour of the speaker that leads his audience to take him seriously when they shouldn't.

irrelevant by deferring to those who are competent in some domain. Sometimes it might be unclear to some whether or not an alternative is relevant, yet extremely clear to competent others because they are experts on the matter at hand or because they occupy the relevant standpoint. It's been convincingly argued that we often have good reason to accept the testimony of those who are experts or who occupy particular standpoints (see, e.g., Wylie 2004, Alcott 2005). The same reasons carry over from how we should respond to testimony about what are appropriate judgements of relevance. We can see this more clearly by adding to **CLASSROOM** that, in addition to the observational evidence possessed by Rashaan, after several years of being repeatedly targeted by racial stereotyping and being sent to the Principal's office for minor offences, he's become attuned to patterns of racial discrimination in his school life. Taking seriously the thought that particular social locations often come with epistemic advantages, when it comes to accessing knowledge about the racism faced by Black people, Black people are well positioned to know when and how it manifests because they experience it frequently and in a patterned way. If this is correct, then this gives us additional reason to take Rashaan seriously when he resists the Principal's suggestion that perhaps he's seeing racism where there is none. Generalising, we can look to those who occupy epistemically privileged standpoints to determine relevance or irrelevance.³⁶

The irrelevant alternative is judged to be relevant by those with epistemic power in the context because of a prejudice they have in favour of the Principal. Both the Principal and the Assistant Principal's epistemic power derives from their institutional as well their social positions. The

³⁶ One might wonder whether there is a tension between my claim here that any alternative that is raised by a speaker who is competent in the relevant domain is relevant, and my rejection of the possibility of letting epistemic authority stand for conversational authority in Blome-Tillmann's contextualism in chapters 4 and 5. But notice that the notion of competence here is different from the notion of epistemic authority discussed earlier. Here, competence amounts to a kind of expertise. The notion of epistemic authority discussed earlier was Zagzebski's (2013) notion of subjective epistemic authority according to which it is a necessary condition on someone being an epistemic authority for you that you trust that they are more epistemically conscientious than you are. So, there is no tension between my rejecting the use of Zagzebski's notion of epistemic authority to develop presuppositional epistemic contextualism and my embrace of the idea that an alternative being raised by someone with relevant competence can make it relevant.

Principal has a positive prejudice towards himself that leads him to overestimate his own ability to detect racism based on his recent training. The Assistant Principal also has a positive prejudice in the Principal's favour that leads him to overestimate the Principal's racism detection capabilities too, and which leads him to take the Principal's alternative seriously. The Assistant Principal thereby confers positive epistemic status on the Principal, enabling the Principal to accumulate even greater epistemic power, and likely boosting the Principal's own self-trust further. The role of the Assistant Principal mirrors the mechanism described by Karen Jones (2012: 245 - 247) by which dominant knowers can scaffold one another's self-trust to excessive degrees and shore up one another's epistemic power beyond what is warranted, thereby creating the conditions in which epistemic injustice can flourish. The Principal is judged not only to be well placed to detect racism but is judged to be better placed than Rashaan. The overestimation claim is plausible. Evidence suggests that the measures employed in anti-racism courses are largely ineffective.³⁷

CLASSROOM involves different harms and consequences from BOARDROOM. Unlike in BOARDROOM, the 'knowledge' of the target of the injustice in CLASSROOM is at issue. This is again connected to the distorting effect of the reception given to the raised alternative on the economy of knowledge ascriptions. Because the Principal's alternative is irrelevant, Rashaan's knowledge ascription remains literally true after the alternative is raised. But because those with epistemic power take the Principal's alternative to be relevant, they mistakenly judge that Rashaan's knowledge ascription is false. This is the distortion: Rashaan is denied knowledge when he shouldn't be because an irrelevant alternative is taken seriously by those with epistemic power when

³⁷ For example, measures that are designed to mitigate the effects of implicit bias have been shown to be ineffective. Lai et al. (2016) have shown that whilst interventions designed to reduce the effect of implicit racial preferences have some effect immediately after the interventions, this effect disappears several hours to several days later.

it shouldn't be. Consequently, Rashaan loses out on the statuses and permissions that come with being ascribed knowledge. He suffers an affront to his epistemic agency in being seen as an unreliable informant. He is perceived as lacking permission to assert and act on the proposition that he was a victim of racism. Being perceived to lack the permission to act, any actions that he takes will be open to criticism, potentially shaking his confidence in the righteousness of his protest and his determination to take action. Thus, alternatives being mistakenly taken to be relevant by those with epistemic power can undermine actions and place barriers in the way to achieving justice.

Like testimonial injustice, this form of irrelevance injustice involves a testifier being unfairly judged not to 'know', but the mechanism that brings about this judgement is quite different. In testimonial injustice a speaker is judged not to know because the audience harbour a prejudice against the speaker. In irrelevance injustice a speaker raises an alternative that is taken seriously when it should not be because of a prejudice in the speaker's favour, which in turn leads an audience to mistakenly judge that an alternative needs ruling out, and to a mistaken knowledge denial.

Looking at BOARDROOM and CLASSROOM side by side reveals an asymmetric ability to influence what's perceived to be a relevant alternative, which in turn affects the ease or difficulty with which the differentially socially situated agents are ascribed knowledge. Those who have their alternatives taken seriously can increase the size of the set of alternatives that others must rule out prior to being ascribed knowledge, whilst the reverse is not true. Consequently, there is an asymmetry in the amount of epistemic labour that differentially situated agents must expend, both in gathering evidence and in demonstrating their possession of such evidence, prior to being ascribed knowledge. Miss Triggs not being taken seriously reduces the amount of epistemic labour that Mr X must expend in order to be ascribed knowledge. The Principal being taken seriously increases the amount of epistemic labour that Rashaan must expend in order to be ascribed

knowledge. Irrelevance injustice shows how the powerful can make epistemic life easier for themselves and harder for others.

Now for a third and final case.

DINING CLUB: Charles, a white man hailing from an elite social class who has just made the transition from being a student at an elite public school to being a student at the University of Oxford, is having dinner at the Oxford Bullingdon club. The topic of conversation is a recent survey report revealing that an astonishingly high number of women have experienced sexual harassment as students at the University. Charles raises the skeptical possibility that many of the reported cases are really instances of women overreacting to harmless flirting. The other diners unanimously and vociferously agree.³⁸

DINING CLUB presents a third case of irrelevance injustice which is importantly different from the other two we've looked at. In **DINING CLUB**, an alternative is taken seriously when it should not be³⁹, but not because of any prejudice toward the speaker. Rather, the alternative is taken seriously because the audience harbour a prejudice pertaining to the alternative's *content*. As a first pass we can say that the audience harbours a content prejudice in that they are prejudiced against a group that the alternative is about. I'll develop a more thoroughgoing definition below. The diners hold the prejudicial belief that women are unreliable in their judgements about whether they have been sexually harassed.⁴⁰ They would have taken the alternative seriously whatever the identity of the speaker. This shows that irrelevance injustice can be a matter of an alternative's content in the sense that it involves an alternative being problematically received

³⁸ The case is fictional, but the belittlement of women has been well documented as a topic of conversation at meetings of the Bullingdon club (see, e.g., Sherwood 2019).

³⁹ A recent meta-analysis found that accusations are 95% reliable (Ferguson & Malouff 2016). Given that the skeptical possibility raised by Charles is very distant from the truth it should be deemed irrelevant.

⁴⁰ This is a common prejudice. Deborah Turkheimer (2017) cites a survey report of close to 900 police officers in which over half of respondents said that between 10 - 15% of sexual assault complainants lies about being assaulted.

because of its content, rather than because of who raises it. This marks another important difference between irrelevance injustice and testimonial injustice. Testimonial injustice, as Fricker (2007) develops it, always involves an identity prejudice towards the speaker. **DINING CLUB** shows that this is not required for irrelevance injustice.

Unlike in **BOARDROOM** and **CLASSROOM**, those who are harmed by the injustice that takes place in **DINING CLUB** are not themselves participants in the conversation. Rather, the reception given to the raised alternative causes distortions in the economy of knowledge ascriptions that harm people outside of the conversation. Particular alternatives being seen as irrelevant in one context can reinforce what alternatives are seen as relevant in other contexts. Imagine that having taken Charles's skeptical possibility seriously, those inside the conversation less readily ascribe knowledge to women they later encounter who claim to have experienced sexual harassment because they apply skeptical standards. The women are seen as lacking the statuses and permissions that come with being ascribed knowledge. The judgement that the women are unreliable informants reinforces the stereotype that women are unreliable testifiers about sexual harassment. They will be judged as not being positioned to act on the proposition that they were sexually harassed. Such judgements may feed into criticisms of actions, and such criticisms may stop the women in their tracks, undermining the pursuit of justice.

With these three examples in mind, we can say that irrelevance injustice occurs when:

- (i) S raises p as an alternative to some proposition already under discussion, q; and
- (ii) p is not taken seriously when it is relevant, *or* p is taken seriously when it is not relevant, by those members of S's audience that possess epistemic power in the context;

- (iii) for the reason that *either* S's audience are prejudiced towards S; *or* S's audience are prejudiced towards the content of what S says; *or* taking/not taking p seriously conforms to a general pattern of oppression.

I will now discuss each of the conditions listed in greater detail. In order to establish its distinctiveness, in section 4 I will spend some more time comparing irrelevance injustice with two other forms of epistemic injustice: testimonial injustice (Fricker 2007) and contextual injustice (Ichikawa 2020).

Condition (i): raising alternatives

Condition (i) has three main components. Firstly, 'proposition already under discussion, q' refers to the proposition that's initially being discussed to which an alternative is raised. The proposition might be under discussion by figuring in a knowledge claim, such as the proposition that Punch's recent decline in profits is explained solely by spending inefficiencies that figures in the self-ascription of knowledge to Mr X in BOARDROOM, or the proposition might come up as part of a discussion of whether it's true, such as the proposition that an astonishingly high number of women have experienced sexual harassment as students in DINING CLUB.

Secondly, an alternative to the proposition under discussion is one that's incompatible with the proposition that's initially under discussion. For example, that the company's decline in profits is (at least partially) due to structural changes in the economy is incompatible with the proposition that Punch's recent decline in profits is explained solely by spending inefficiencies. Whether we attribute knowledge depends on whether we judge that the subject can rule out the alternatives that

are taken to be relevant. Whether knowledge claims are in fact true depends on whether the subject of the knowledge ascription can in fact rule out alternatives that are in fact relevant.

What does raising an alternative – the third component of condition (i) – amount to? One possible answer conceives of raising an alternative as a distinct speech act, one that has its own set of felicity conditions that must be satisfied for its successful performance. But, a less theoretical account will do for our purposes: a speaker raises an alternative whenever they bring an alternative to salience. This can be achieved with different sentence types. For example, I might raise an alternative by uttering a declarative sentence, as Miss Triggs does when she declares that ‘the economy has undergone structural changes’ in BOARDROOM. Or I might raise an alternative by asking the question ‘but what about not-p?’, as the Assistant Principal does when he asks of Rashaan ‘Can you rule out the possibility that you’ve misinterpreted Ms. H’s actions?’ in CLASSROOM. Or I might bring an alternative to salience by asserting a proposition that is incompatible with q.

Condition (ii): taking alternatives seriously

Understanding condition (ii) requires: an account of when an alternative is, and is not, taken seriously; an account of when an alternative should, and should not, be taken seriously, and; a sketch of what it takes to possess power in a context. Let’s take each in turn.

Let’s say, following Blome-Tillmann (2009, 2014), that a speaker takes an alternative seriously when that alternative is compatible with their pragmatic presuppositions:

Pragmatic Presupposition: S pragmatically presupposes p in C iff S is disposed to behave, in her use of language, as if she believed p to be common ground in C (Blome-Tillmann 2014: 26, amended from Stalnaker 1974).

Where common ground is:

Common Ground: It is common ground that *p* in a group *G* iff all members of *G* accept (for the purposes of the conversation) that *p*, and all believe that all accept that *p*, and all believe that all believe that all accept that *p*, and so on (Stalnaker 2002: 704).

Applied to the cases we've been discussing, in **BOARDROOM** all conversational participants, apart from Miss Triggs, do not take the proposition that Punch's decline in profits is due to structural economic changes seriously because they are disposed to behave as if they believe the incompatible spending inefficiency hypothesis to be common ground. Miss Triggs raises the alternative, and all other conversational participants continue to ascribe knowledge to Mr X as if Miss Triggs had never spoken. In **CLASSROOM** the Principal and his Assistant take seriously the proposition that Rashaan sees racism where there is none in that it is compatible with what they are disposed to behave as if they believe is common ground. They act as if they believe Rashaan does not 'know' that he suffered from the racial prejudice of his teacher, and it's compatible with this that Rashaan misinterpreted the actions of Ms. H. Finally, in **DINING CLUB** Charles's audience take seriously the proposition that many cases that women report as sexual harassment are really overreactions to harmless flirting in that they are disposed to behave as if it's common ground.⁴¹

⁴¹ Making use of the machinery of speech-act theory and following debates about silencing, we might instead conceive of taking and failing to take seriously as a matter of secured and failed uptake. There are different ways to think about uptake. Traditionally uptake is thought to be active in that it is secured when an audience recognises the intention of the speaker (Austin 1962: 117). More recently, Langton has pointed out that uptake can also be passive (2018). Langton's model of passive uptake invokes the phenomenon of accommodation: the process of constant adjustment to the conversational score that makes a move in a conversation count as 'correct play'. Langton points out that the accommodation of some speech acts requires the accommodation of particular presuppositions. Presupposition accommodation does not require any active state in a hearer's mind, contrary to active uptake. Rather, it can occur by the mere omission of hearers. However, Langton points out, sometimes presuppositions are not allowed to pass. Instead, they are 'blocked'. When hearers block what is presupposed, they at the same time block the speech act to which the presupposition contributes. In this way, blocking can lead to uptake failure, illocutionary disablement, and silencing.

We can make use of this machinery to provide an alternative model of how irrelevance injustice works. The speech act of raising of an alternative requires that the audience presuppose that the speaker has the relevant credibility and/or authority to successfully pull off raising an alternative. But, in **BOARDROOM** the presupposition that Miss Triggs has

An alternative should be taken seriously when it is relevant, and should not be taken seriously when it is not relevant. I made some suggestions about what makes an alternative relevant in the course of presenting the cases of irrelevance injustice. On my view two ways that speakers can make alternatives relevant is by raising them in an act of intellectual virtue – a view that I will develop further in chapter 7 – and by raising them when competent in the relevant domain. An alternative is irrelevant when it is not ruled in by any of the rules of relevance. I don't intend these considerations to be exhaustive of what determines relevance. In light of Kristie Dotson's (2012: 41 – 42) caution against closed conceptual structures that risk inadvertently perpetuating new forms of epistemic injustice at the same time as theorising about it, it'd be unwise to claim to have settled when an alternative is and is not relevant, and thus when we have an instance of irrelevance injustice on our hands. Note also that the definition of irrelevance injustice is compatible with different accounts of what makes an alternative relevant. This means that even if you don't agree with my proposals for what makes alternatives relevant, you can still accept the reality and importance of the phenomenon of irrelevance injustice.

An interesting payoff of developing a theory of irrelevance injustice is that it can help us to establish further rules of relevance. We might encounter cases that we clearly judge to be cases of irrelevance injustice. We might then use our judgements about such cases to develop an account of what makes an alternative relevant. For example, the theory of irrelevance injustice provides us with a new

the credibility and/or authority required to successfully raise an alternative is blocked, and Miss Triggs is thereby silenced. Applied to CLASSROOM, Rashaan simply does not have the power that is required to block the presupposition that his Principal has the requisite credibility and/or authority to raise the alternative that he does. This enables the Principal to assume credibility and/or authority on matters of race that he simply does not have, and undermine Rashaan's claim to knowledge in the process.

There are good reasons to prefer my epistemic injustice model of irrelevance injustice over a speech-act model. On a silencing model uptake is a necessary condition for a speaker successfully raising an alternative. This means that in BOARDROOM-type cases, because the target of the injustice cannot be said to have successfully raised an alternative, the perpetrator of the injustice cannot be faulted on the grounds of not taking the speaker's alternative seriously. But clearly the perpetrator is criticisable. This suggests that my epistemic injustice model is the more empirically accurate description of what occurs in such cases.

reason to reject Lewis's rule of attention. It says that any possibility that is attended to is relevant (1996: 559). Accepting the rule of attention would lead us to fail to diagnose cases of irrelevance injustice, and this is a reason to reject it. The rule of attention entails that as soon as any alternative is mentioned it becomes relevant, and so would dictate that there is no injustice in **CLASSROOM** or **DINING CLUB** given that the Principal and Charles mention the problematic alternatives.

The final component of condition (ii) is an audience's possession of epistemic power in the context. This component was unpacked in the discussion of the central cases.

Condition (iii): patterns of prejudice

Let's work through the three disjuncts of condition (iii). The first concerns prejudice towards the speaker. On Fricker's account, testimonial injustice occurs when a speaker is not given the credibility they deserve when testifying because of the operation of negative identity-prejudicial stereotypes in the audience. The notion of a negative identity-prejudicial stereotype needs unpacking. On Fricker's account an identity-prejudice is a prejudice against or in favour of an individual in virtue of their membership of some social group. Prejudices can be felt through stereotypes, which Fricker takes to be widely held associations between a social group and one or more attributes (2007: 35). Where identity-prejudicial stereotypes are negative, the association involved is a disparaging one. Identity-prejudicial stereotypes are resistant to counterevidence. Negative identity-prejudicial stereotypes, then, are disparaging stereotypes employed in our judgements of people that are resistant to counterevidence. Testimonial injustice occurs when such stereotypes are employed in our judgments of a person's testimony such that we're led to give them deflated credibility when testifying. Applying the notion of negative-identity prejudicial stereotypes to irrelevance injustice, we can say that irrelevance injustice occurs when a relevant alternative is

raised, yet the alternative is not taken seriously by the audience because of the operation of a negative-identity prejudicial stereotype. In **BOARDROOM** Miss Triggs suffers from irrelevance injustice in that the other conversational participants wrongly fail to take her alternative seriously given that they harbour a negative-identity prejudicial stereotype that says that women are not well placed to make claims about economics and business.

We must depart from Fricker by allowing for positive identity-prejudicial stereotypes to figure in irrelevance injustice too. Fricker rejects the possibility of positive identity-prejudicial stereotypes figuring in testimonial injustice because she thinks that inflated credibility assessments don't generate harms (Fricker 2007: 21).⁴² Following Fricker here would prevent us from making sense of cases involving the taking seriously of alternatives when they should not be. For example, in **CLASSROOM** it's the overly inflated judgement of the Principal's ability to discern when racism exists that leads both the Principal and his Assistant to take the raised alternative seriously. This explains why Rashaan is wrongfully denied knowledge and the statuses and permissions that come with it.

The second disjunct also concerns prejudice, but the prejudice attaches to what the speaker says rather than to the speaker's identity. While discussing **DINING CLUB I** I glossed content prejudice as 'prejudice against a group that the alternative is about'. While this is surely one way to harbour content prejudice, it will not do as a definition. It is both too broad and too narrow. It is too broad because it implies that an audience has a content prejudice toward every alternative that is about a group against whom they harbour some prejudice. Charles is prejudiced against women, but he is surely not prejudiced against every alternative that is about women. It is too narrow because it only

⁴² See Medina (2011) for a rejection of this point and for an argument that testimonial injustice can follow from credibility excess.

allows for prejudice toward one class of speech act, namely, alternatives. Robin Dembroff and Dennis Whitcomb (forthcoming) have recently argued that there is a class of epistemic injustices that they call content-focused injustice which is ‘epistemic injustice focused not on the identity of a given speaker, but rather on the things that the speaker communicates’ (forthcoming: 3). Dembroff and Whitcomb discuss cases in which audiences have prejudices against a group that some assertion is about or associated with (forthcoming: 7). **DINING CLUB** shows that irrelevance injustice can itself take the form of content-focused injustice. It will be useful to have a definition of content prejudice that illuminates these different varieties of content-focused injustice. Dembroff and Whitcomb do not provide one, rather they elucidate the notion by providing some examples of content-focused injustice. Here’s a definition that I think gets things right:

Content prejudice: An audience A is prejudiced against a content C iff S has an identity-prejudice towards the group that C is about or associated with and A’s identity-prejudice influences A’s assessment of C.

The definition avoids being too broad since it doesn’t imply that Charles is prejudiced against every content that is about women, rather it implies that he has a content prejudice only when his prejudice against women is influencing his assessment of a content. It also avoids being too narrow since it allows for prejudice against any speech act and can thus be employed to illuminate different varieties of content-focused injustice, including those discussed by Dembroff and Whitcomb.

The third disjunct follows a suggestion made by Derek Anderson (2017: 211). The judgement that a raised alternative is not relevant when it is, or is relevant when it is not, does not have to be caused by prejudice for it to count as an injustice. Rather, what’s important might be whether the judgement of irrelevance or relevance conforms to a general pattern of privilege and oppression. This disjunct was motivated by the variant of **BOARDROOM** that I described above, in which the men don’t

take Miss Triggs's alternative seriously because they're threatened by her. While this case doesn't involve prejudice, it surely involves misogyny (Manne 2017). I want to provide another variant of **BOARDROOM** in which the perpetrators do not harbour any malicious attitudes in order to illustrate that injustice can be perpetrated and systems of oppression can be maintained in the absence of any malicious individual attitudes. Suppose that the reason that the men in **BOARDROOM** fail to take the alternative that Miss Triggs raises seriously is not because they are prejudiced against women, but rather because they are committed to an economic theory that implies that the considerations that Miss Triggs raises can only have had a negligible effect on Punch's profit margins. Suppose that men tend to believe this theory more frequently than women, and that this commitment regularly gives rise to a failure to take seriously relevant alternatives that are raised by women in business contexts. Making prejudice a necessary condition means this wouldn't count as irrelevance injustice: the wrong result. We can get the right result by making prejudice sufficient, but not necessary. Instances where there is no prejudice, but mere conformity to a general pattern of privilege and oppression, still unfairly privilege some and harm others in the same ways as those judgements that are caused by prejudice. When judgements of relevance conform to the general pattern of norms that work to harm the oppressed by making it more difficult for them to be ascribed knowledge, and privilege the dominant by making it easier for them to be ascribed knowledge, they are of concern even when conformity does not arise from prejudice.

4. Irrelevance Injustice and other forms of epistemic injustice

I'll finish up by noting some important differences and connections between irrelevance injustice and other forms of epistemic injustice in order to dispel any doubts about the distinctiveness of irrelevance injustice.

4.1. Irrelevance Injustice and Testimonial Injustice

Some important differences between irrelevance and testimonial injustice have already been noted. I'll finish up by mentioning three further differences. Firstly, irrelevance injustice and testimonial injustice *don't always co-occur*. BOARDROOM and DINING CLUB-type cases of irrelevance injustice are fundamentally different in structure from testimonial injustice. Whereas in testimonial injustice the victim of injustice is always the person testifying, in these forms of irrelevance injustice the person raising the alternative is the perpetrator. BOARDROOM-type cases of irrelevance injustice are more similar in structure to testimonial injustice, but there will not always be co-occurrence in these cases either. Testimonial injustice arises in a hearer's evaluation of whether a speaker's utterance is known or credible, irrelevance injustice arises in an audience's evaluation of whether an utterance is relevant. This means that there can be cases in which a hearer correctly judges a speaker's testimony to be known, and so does not perpetrate testimonial injustice, yet incorrectly judges the same piece of testimony to be irrelevant, thereby perpetrating irrelevance injustice. Suppose that in BOARDROOM Miss Triggs is correctly judged to know that there have been structural changes in the economy, and thus does not suffer from testimonial injustice. However, what Miss Triggs says is incorrectly judged to be irrelevant to explaining the cause of Punch's losses. Thus, we have a case of irrelevance injustice that is not at the same time a case of testimonial injustice.

Secondly, testifying and raising an alternative as relevant are *distinctive epistemic activities*. One way to see this is to think about how best to classify testimonial and irrelevance injustice in terms of the distinction that Christopher Hookway (2010) draws between epistemic injustices that are perpetrated from the *informational perspective* and those perpetrated from the *participant perspective*. Informational perspective epistemic injustice involves an agent being wrongly treated

as an unreliable source of knowledge. Participant perspective epistemic injustice involves an agent being wrongly taken to lack the competency required to participate in some epistemic activity besides the possession or transmission of knowledge. To characterise irrelevance injustice in terms of this distinction we have to expand it to also include in the category of participant perspective epistemic injustice cases where agents are wrongly taken to *possess* the competency required to participate in epistemic activities besides knowledge transmission. This is necessary to be able to couch BOARDROOM and CLASSROOM type cases in terms of this distinction. With this broadening in mind, we can say that testimonial injustice is an epistemic injustice of the informational perspective variety. As the label has it, it involves the speech act of testifying. The primary epistemic function of testifying is to put forward some proposition as true, and thereby transmit knowledge to an audience. Irrelevance injustice, on the other hand, is an epistemic injustice of the participant perspective variety. It involves the raising of alternatives, the primary epistemic function of which is to inquire into whether a proposition under discussion is true or justified in being believed, and thus whether it's known.

Plausibly, the different epistemic activities involved in testimonial and irrelevance injustice each require the cultivation of different epistemic virtues. Testimony might be associated with the virtue of intellectual courage because courage is required to put forward propositions as something that others can take your word for. When you testify you put forward a proposition as true, and so take the matter to be settled. Raising alternatives is to take the matter to still be open. So, whilst requiring courage, raising alternatives also involves the virtues of open-mindedness and intellectual curiosity. Open-mindedness was discussed earlier. The intellectually curious person is motivated to engage in inquiry, to ask and answer questions, and to improve their epistemic standing as they consider and rule out more alternatives (see, e.g., Ross 2020 on the virtue of curiosity). Thus, suffering from each may lead to inhibiting the cultivation of, and eroding, distinct epistemic capacities. What's

more, lumping together epistemic injustices that involve distinct epistemic activities may obscure the differences in the solutions required to address each. Fricker's 'virtue of testimonial justice' might work as a response to testimonial injustice (though see, e.g., Anderson (2012), Alfano (2015) for a critique), but not to irrelevance injustice. Treating the two as the same phenomena can obscure this fact.

One final point worth registering is that testimonial and irrelevance injustice arise from an audience's failure to comply with different norms. Testimonial injustice arises when a hearer fails to comply with an evidentialist norm of credibility because of the operation of an identity prejudice, where the evidentialist norm of credibility says that a hearer, *H*, should match their credibility judgement of *S* to the evidence that *S* is offering the truth (Fricker 2007). Irrelevance injustice, by contrast, arises when a raised alternative is either wrongly not taken seriously or wrongly taken seriously as dictated by the rules of relevance. This is important to note because it shows that if one is to avoid the perpetration of both testimonial and irrelevance injustice, one must acknowledge and comply with two different sets of norms. Merely complying with one will be insufficient.

4.2. Irrelevance Injustice and Contextual Injustice

Contextual injustice, as applied to knowledge ascriptions, occurs when the standards for 'knows' are prejudicially raised to an inappropriate level, and a testifier is wrongfully denied to 'know' as a result (2020: 16). In contextual injustices, raising alternatives to salience does in fact raise the standards such that the knowledge denial is literally true, but, Ichikawa claims, it is nevertheless unjust because 'inappropriate' (2020: 9).⁴³

⁴³ Ichikawa develops his notion of contextual injustice on the standards model of contextualism rather than the relevant alternatives model.

Irrelevance injustice and contextual injustice are not entirely co-extensive. For one thing, contextual injustice is always a matter of wrongful knowledge denials. But in **BOARDROOM**-type cases of irrelevance injustice whether the victim 'knows' is never at issue. For another thing, Ichikawa seems to make negative identity-prejudice against the testifier or against the content of the testimony a necessary condition for contextual injustice (2020: 16). For the reasons described in the discussion of condition (iii) of the definition of irrelevance injustice, I think this is too restrictive. We will miss out on classifying cases in which an irrelevant alternative is taken seriously when it should not be because of positive identity prejudice – such as in **CLASSROOM** – and cases in which judgements of irrelevance or relevance conform to a general pattern of bias against marginalized groups or in favour of dominant groups – such as in modified **BOARDROOM**.

Irrelevance injustice and contextual injustice are not entirely disjoint either. Consider a version of **CLASSROOM** in which the reason that the Principal raises the skeptical alternative is because of racist prejudice that black people often see racism where there is none. This version of **CLASSROOM** counts both as a case of irrelevance injustice and as a case of contextual injustice. But the explanation for why it's an injustice is different on each theory. On Ichikawa's view, contextual injustice involves the literal manipulation of the truth-conditions of knowledge ascriptions. When the Principal raises the skeptical alternative it is made relevant. Because Rashaan can't rule it out, his knowledge ascription is false, and the Principal speaks truly when he denies that Rashaan 'knows'. The injustice, on Ichikawa's view, is that the standards for 'knows' that the Principal invokes are unfair. So, while the Principal's knowledge denial is literally true, it is nevertheless unjust. On my view, by contrast, the skeptical alternative that the Principal raises is irrelevant, and so does not alter the truth-conditions of the knowledge ascription. Yet, the alternative is perceived to be relevant by those who possess epistemic power in the context, and this leads them to mistakenly deny that Rashaan 'knows'. While the Principal's denial of Rashaan's knowledge is

literally false, it is perceived as being true by those with epistemic power in the context. Rashaan suffers an injustice because knowledge is denied where it should be ascribed.

Which analysis should we prefer? The answer comes down to whether or not the skeptical alternative is in fact relevant. Ichikawa says that it is, I say that it is not. In other work Ichikawa makes it clear that he rejects both Lewis's rule of attention and Blome-Tillmann's rule of pragmatic presupposition. He writes: 'Lewis's Rule of Attention is too weak' (2017: 25). In (Ichikawa 2015) he proposes a series of counterexamples to Blome-Tillmann's proposal that he takes to be grounds for rejecting the view. Why then, on Ichikawa's view, is the Principal's alternative relevant? It's not clear. Rather than trying to figure out when exactly Ichikawa thinks an alternative is relevant and then trying to settle who's right about relevance - a task that has occupied contextualists for decades - I want to point to an area of agreement that gives us reasons for having both the concept of irrelevance injustice and the concept of contextual injustice. One thing that we both agree on is that cases like CLASSROOM involve an injustice and that we need concepts that help us to track the injustice. We disagree about the explanation. Intuitions vary. I suspect that some will agree with me about the relevance of the Principal's alternative, and that others will agree with Ichikawa. The fact that intuitions vary and that there doesn't seem to be any way to settle the dispute about relevance anytime soon justifies having both concepts. The goal of tracking the injustice in cases like CLASSROOM is best served by having both concepts. If your intuitions are with Ichikawa then you have a case of contextual injustice. If your intuitions are with me, then you have a case of irrelevance injustice. Crucially, wherever your intuitions lie, you have managed to track the injustice, which is the reason for developing such theories in the first place. There's an injustice for both intuitions. Thus, the fact that the two concepts don't have entirely different extensions doesn't undermine the importance and usefulness of having both of the concepts (see Anderson 2018 for

further discussion of the point that two concepts don't have to have entirely different extensions for it to be worth having both).

5. Conclusion

This chapter has delivered the second payoff of theorising about epistemic contextualism non-ideally. Investigating the relationship between social context, conversational context, and the mechanisms of raising alternatives to relevance emphasized by epistemic contextualists has shown up an unrecognized form of epistemic injustice: irrelevance injustice. Irrelevance injustice occurs either when a speaker advances an alternative that is not taken seriously when it should be, or when a speaker advances an alternative that is taken seriously when it should not be. Irrelevance injustice causes distortions in the economy of knowledge ascriptions, affecting how much epistemic labour members of different social groups must undertake in ruling out alternatives prior to being ascribed knowledge, and thus whether they end up being ascribed knowledge. Given that knowledge ascriptions have socially valuable functions, irrelevance injustice generates asymmetries in who gets to take advantage of these socially valuable functions.

Chapter 6.5: Taking Stock of Conversational Contextualism's Problems

As a result of theorising contextualism non-ideally by removing the common contextualist assumption that speakers have equal abilities to shift the standards for 'knows' and replacing it with the more realistic non-ideal assumption that speakers have unequal abilities to shift standards that varies with unfairly distributed power, we've seen the problems for conversational contextualism mount up. We've seen new classes of counterexamples, we've seen morally and politically troubling consequences that stem from the interaction of contextualism, power, and the knowledge norms, and we've seen an inability of contextualism to diagnose epistemic injustice. We've seen that some of these problems cause trouble for specific versions of contextualism, and that others cause trouble for conversational contextualist views more broadly. We've seen that all of these problems share a common core: conversational contextualists tie the truth-conditions of knowledge ascriptions to conversational dynamics without placing adequate constraints on the effects of conversational dynamics on the truth-conditions of knowledge ascriptions. How should conversational contextualists respond to these problems? There are at least three options: (1) dump contextualism once and for all; (2) attempt to save conversational contextualism by finding a way to appropriately constrain the effects of socially situated and power infused conversational dynamics on the truth-conditions of knowledge ascriptions, or; (3) develop a non-conversational contextualism that ties the truth-conditions of knowledge ascriptions to some feature of the context other than conversational dynamics.

I suspect that at this stage in the thesis some readers will be thinking that option 1 is surely the only way to go. After all, the way that this thesis has panned out has made it seem as though contextualism brings us nothing but trouble. But this is to overlook the fact that a workable contextualism promises

significant philosophical payoffs. Why else would philosophers have been trying to make the view work for over 50 years? (See DeRose 2009 chapter 1 & Rysiew 2021 for helpful histories of contextualism). Contextualism promises a resolution to perennial epistemological puzzles like skepticism (see, e.g., Cohen 1986, 1988, 1999, DeRose 1995, Heller 1999, Lewis 1996, Neta 2004, and Stine 1976 for contextualist resolutions to skepticism), and the lottery paradox (see, e.g., Cohen 1988 1998, Lewis 1996, Neta 2002, Reiber 1998 for contextualist resolutions to the lottery paradox). It can explain patterns of linguistic intuitions, such as about DeRose's (1992) bank case, and Cohen's (1999) airport cases. And contextualism can be recruited in the service of highlighting unjust aspects of our socio-epistemic landscape (see, e.g., chapter 6 of this thesis, and Ichikawa 2020). Given that contextualism promises so much, it's worth exploring all the options to make it work.

The final chapters of this thesis pursue options (2) and (3), and in so doing deliver on the third payoff from theorising contextualism non-ideally: we get two new forms of contextualism that are motivated by the problems identified in the earlier parts of the thesis. Chapter 7 pursues option (2). It starts out with the observation that all of the problem cases identified in chapters 3, 4, and 5 involve nefarious actors raising alternatives that undermine the truth of intuitively plausible knowledge claims. Taking this observation as its motivation, chapter 7 develops a virtue epistemology inspired contextualism that excludes the conversational moves of nefarious actors from contributing to the truth-conditions of knowledge claims. The proposed view can properly be called a version of conversational contextualism because it makes the truth-conditions of knowledge claims dependent on the dynamics of the conversation in which the claim is made, albeit in a highly constrained way. Chapter 8 pursues option (3). It starts out with the observation that all the problem cases identified in chapters 3, 4, and 5 of the thesis share another commonality: they all involve high moral stakes. Taking this observation as its motivation, chapter 8 develops a version of stakes

contextualism that divorces the truth-conditions of knowledge claims from the conversational context altogether, and instead ties them to what's morally at stake in the context. This builds on existing versions of stakes contextualism that take practical stakes to affect the truth-conditions of knowledge claims and suggests that various kinds of moral stakes can also have an effect. I stop short of endorsing either of these new forms of contextualism outright, and instead rest content with showing that each can handle the problems identified for existing forms of contextualism earlier on in the thesis whilst retaining some of their key selling points, and that as such both are worthy of further development.

But just before we get to that, it'll be worth refreshing our memories of the problems that we're looking to overcome in developing new contextualism.

6.5.1. Chapter 3: Scorekeeping Trolls

In Chapter 3 I argued that Keith DeRose's (2004, 2009) gap view is vulnerable to trolling. I argued that trolling results in numerous counterintuitive implications which taken together constitute a reductio of the gap view. The structure of trolling cases are as follows: at t1 a knowledge ascription (denial) is true/false in a conversation because it meets/fails to meet each of the participant's epistemic standards. Later, at t2, a troll enters the conversation with different standards - either higher, such as in VACCINE, or lower, such as in BRIDGE - which are such that any future tokens of the same knowledge ascription (denial) type meet/fail to meet, these standards. This means that future tokens of the knowledge ascription (denial) type will be truth-valueless. This counter-intuitive result occurs regardless of how outlandish or inappropriate the troll's standard appear to be. It's because DeRose requires that the personally indicated standard of every speaker register in the truth-conditions for knowledge ascriptions that the gap view has this implication.

The counterintuitive implications of DeRose's gap view are compounded when the problematic verdicts it reaches on trolling cases are considered in conjunction with knowledge norms of **ASSERTION**, **BELIEF**, and **ACTION**. These norms say that it is permissible for *S* to assert, believe, and act on *p* iff *S* satisfies 'knows that *p*'. In trolling cases 'knows that *p*' is rendered truth-valueless, and so is not satisfied. The upshot is that the trolled party cannot permissibly assert, believe, or act on *p* merely because a troll has entered the conversation with different epistemic standards from existing participants. This is counterintuitive: whether or not you're epistemically permitted to assert, believe, or act on *p* should not so precariously depend on who you are having a conversation with. We'll be looking for a version of contextualism that prevents trolls from undermining people's intuitively true knowledge claims.

6.5.2. Chapter 4: Power and Presuppositional Epistemic Contextualism

Chapter 4 provided an argument against Michael Blome-Tillmann's (2009, 2014) presuppositional epistemic contextualism (PEC). I identified a host of problematic implications that follow from combining PEC with real world power relations, and the knowledge norms, and argued that taken together these implications are grounds for rejecting PEC.

We saw that because what's pragmatically presupposed in a conversation can be influenced by power, PEC is vulnerable to two classes of counterexample. The first class of counterexamples were authority problem cases. These are cases where PEC delivers the verdict that an intuitively true knowledge ascription is false because conversational authority is wielded to ensure that an ineliminable skeptical hypothesis is mutually pragmatically presupposed. It is the role that Blome-Tillmann gives to conversational authority that makes PEC vulnerable to authority problem cases. In order to avoid 'knows' being too hard to satisfy, PEC requires that the negation of skeptical

hypotheses are mutually pragmatically presupposed. When pragmatic presuppositions diverge PEC fails to deliver the intuitively correct verdict, and we instead end up in a ‘defective context’. To deal with this, Blome-Tillmann argues that a conversational authority can put their foot down to settle what’s to be pragmatically presupposed thereby ensuring that the negation of skeptical hypotheses are mutually pragmatically presupposed. But because conversational authority can also be wielded to insist on sceptical hypotheses that undermine intuitively true knowledge ascriptions conversational authority becomes source of counterexamples. Our principal authority problem case is **SEXUAL HARASSMENT***.

The second class of counterexamples were adaptive presupposition cases. These are cases where PEC delivers the verdict that an intuitively true knowledge ascription is false because someone alters what’s mutually pragmatically presupposed in a skeptical direction, where the success of altering pragmatic presuppositions depends upon an unjust distribution of power. The feature of PEC that gives rise to adaptive presupposition cases the central role given to pragmatic presuppositions in the theory. Because what one pragmatically presupposes in a conversation can be a function of unjust power relations, socially powerless individuals can fail to challenge sceptical hypotheses raised by more powerful individuals because they’re fearful of the consequences. When powerless individuals are unable to eliminate such unchallenged sceptical possibilities, their intuitively true knowledge ascriptions come out as false according to PEC. Our principal adaptive presupposition case is **SEXUAL HARASSMENT****.

The counterintuitive verdicts that PEC delivers in authority problem cases and adaptive presupposition cases are compounded when the verdicts are considered in conjunction with the knowledge norms. The verdicts that PEC delivers, and the knowledge norms taken together mean that the subject of the intuitively true knowledge ascriptions in the cases are not permitted to assert,

believe, or act on the proposition that they claim to ‘know’. This is counterintuitive: given that we judge that the subject does ‘know’ in these cases, we’ll also judge that they’re not violating the knowledge norms when they assert, believe, or act on the proposition that they claim to ‘know’. And, as our principal cases showed, this can be morally and politically problematic. Taken together this set of counterintuitive implications identified in Chapter 4 are grounds for rejecting PEC. With all that in mind, we’ll be looking for a version of contextualism that avoids authority problem and adaptive presupposition cases, and the politically problem implications that arise when the cases are considered in conjunction with the knowledge norms.

6.5.3. Chapter 5: Epistemic Contextualism or Epistemic Injustice?

In chapter 5 I argued that extant versions of conversational contextualism cannot deliver the verdict that certain clear cases of testimonial injustice that take the form of TI-‘Knowledge’ are in fact cases of testimonial injustice. Consequently, I argued, contextualists face a dilemma: either reject conversational contextualism in its current forms in order to vindicate the intuition that such cases are in fact instances of testimonial injustice, or concede that such cases are not in fact instances of testimonial injustice. We’ll be looking for a version of contextualism that vindicates our testimonial injustice intuitions about cases like **DISMISSED FROM CLASS**, and so which avoids the dilemma that’s faced by extant versions of conversational contextualism.

With the problems fresh in our minds, let’s see if we can develop a new contextualism that solves them.

Chapter 7: Virtue Responsibilist Contextualism

1. Introduction

One thing that unifies all of the cases that we've identified as causing trouble for conversational contextualism is that they all involve a speaker raising an alternative for nefarious reasons: trolling cases involve trolls raising alternatives to derail the conversations and goals of their interlocutors. Our principal authority problem case involves the wielding of conversational authority to raise an alternative for a self-serving end. Our principal adaptive presupposition case involves someone raising an alternative because of their misogynistic attitudes. 'No testimonial injustice' cases involve an alternative being raised for reasons of prejudice. This commonality is instructive in helping to figure out how to adequately constrain the effects of conversational dynamics on the truth-conditions of knowledge claims: we can prevent nefariously motivated conversational moves registering by only permitting speakers' alternatives as relevant when they're raised with the right kind of motivation. The right kind of motivation is epistemic: when the alternative is raised because of a desire to attain an epistemic good like truth or understanding. This epistemic motivation is exhibited when a speaker raises an alternative in an act of intellectual virtue. Making use of these ideas, in this chapter I'll propose a novel version of conversational contextualism that I call *virtue responsibilist contextualism*. As I develop it, virtue responsibilist contextualism is a version of Lewisian relevant alternatives conversational contextualism which draws on resources from virtue epistemology to develop a replacement for Lewis's problematic rule of attention (see chapters 4, 5 and 6 for discussion of why the rule of attention is problematic). According to virtue responsibilist contextualism, if a speaker raises an alternative in an act of intellectual virtue, then it is relevant. None of the alternatives that are raised in our problem cases are ruled in by this rule, nor are they ruled in by any other of Lewis's rules of relevance. This means that virtue responsibilist contextualism delivers the right verdicts on all our problem cases.

What's more, I'll argue that virtue responsibilist contextualism retains a number of the benefits of the extant versions of conversational contextualism that have been the subject of criticism earlier in this thesis. I won't be arguing that virtue responsibilism is a difficulty-free view. One issue that I'll confront later is that because virtue responsibilist contextualism ties the relevance of alternatives to speakers' motivational states, it will sometimes be difficult to tell whether or not an alternative is relevant. Rather, what I'll be arguing is that because virtue responsibilist contextualism both deals our problem cases and retains some key selling points of extant versions of conversational contextualism, it is a view that is worthy of further development.

I proceed as follows: across sections 2 and 3 I consider and dismiss a previous attempt to develop a virtue epistemology inspired version of epistemic contextualism, namely, John Greco's (2004, 2008) virtue contextualism. In section 4 I propose an alternative virtue epistemology inspired version of epistemic contextualism, namely, virtue responsibilist contextualism. I'll show that virtue responsibilist contextualism deals with the problem cases that I've shown spell trouble for extant versions of conversational contextualism. In section 5 I argue that virtue responsibilist contextualism retains some of the key benefits of extant versions of conversational contextualism. In section 6 I consider an objection to virtue responsibilist contextualism. Section 7 concludes.

2. Virtue Contextualism

There are two broad camps in virtue epistemology: virtue responsibilists and virtue reliabilists. Both camps agree that intellectual virtues are cognitive excellences, but they disagree about what these cognitive excellences consist of. Virtue responsibilists (e.g., Zagzebski 1999, Montmarquet 1993) follow Aristotle in claiming that virtues are intellectual states of character that are (i) acquired, (ii) stable, (iii) enable their possessor to reliably realize a certain end⁴⁴, and (iv) include

⁴⁴ Though, there is some disagreement as to whether this condition is necessary. See Battaly (2008: 645) for discussion on this point.

a motivation that is suitable to realize the end in (iii). Virtue responsibilists take paradigm intellectual virtues to include open-mindedness, intellectual autonomy, and intellectual courage. Virtue reliabilists (e.g., Sosa 1991, Greco 2002) follow Plato in claiming that virtues are qualities of a thing that enable it to play its function well. Since our primary intellectual function is attaining truth, the intellectual virtues are whatever faculties enable us to attain the truth. Virtue reliabilists thus reject (i) and (iv) of the responsibilist conception. Instead, virtues are taken to be faculties that reliably enable us to get at the truth whether they be natural or acquired. Reliabilists take paradigm intellectual virtues to include reliable vision, memory, or good reasoning.

Numerous virtue epistemologists claim that knowledge is true belief resulting from intellectual virtue (e.g., Sosa 1991, Zagzebski 1999, Riggs 2002, Lehrer 2000, Greco 2004). Different virtue epistemologists cash this thesis out in slightly different ways depending on what they think the intellectual virtues are. Greco states the thesis in a way that reflects his virtue reliabilism:

VK: In cases of knowledge, S believes the truth because her belief is produced by an intellectual excellence (a cognitive ability or power); for example, S reasoned well, or S remembered correctly, or S saw accurately (Greco 2004: 384).

Greco (2004; 2008) argues that VK entails a form of contextualism about knowledge ascriptions that he calls virtue contextualism. He reasons as follows:

1. Knowledge attributions involve causal explanations.
2. Causal explanations require a context-sensitive semantics.
3. Therefore, knowledge attributions require a contextualist semantics (Greco 2008: 421).

The first premise is just a feature of VK. VK has it that knowledge ascriptions involve ascriptions of causal responsibility: to say that ‘S knows that p’ is to say that S has a true belief that is caused by S’s intellectual virtue.

To get a grip on the second premise, let's consider how it is that causal explanations are context sensitive. When we say that A occurred because B occurred, we're not saying that B is the sole factor that brings about A; rather, we're saying that of all the factors that causally contributed to A occurring, B is a particularly salient causal factor. One determiner of the salience of causal factors is abnormality, and what is abnormal in a given case will be a matter of context. To use Greco's own example (2004: 385), if an elephant walks in the room and causes a panic, we're likely to say that the elephant walking into the room caused the panic. But the elephant walking in is not by itself sufficient to cause the panic. It would not cause a panic if the room were a zoo and the people in the room were zookeepers awaiting the elephant's arrival. But if the room is a pub where all of the customers have the standard reaction to wild animals, then the elephant walking in will be said to be the cause of the panic. Greco describes abnormality as being a feature of the 'subject context': it is a matter of where the event that a causal explanation is being offered for takes place.

Further determiners of the salience of causal factors are the interests and purposes of those participating in the conversation in which the causal explanation is offered. Our interests and purposes will mean that we often pick out certain causal factors as salient. For example, we might pick out factors over which we have control. If a student comes to me and asks why they did so poorly on their essay, given that their interests and purposes are to improve, I'm likely to pick out things that they have control over as causally salient, such as the fact that the essay was poorly structured. And I'm likely not pick out things over which they have no control, such as the fact that university policy dictates that essays are to be marked on a curve. Greco describes interests and purposes as being a feature of the 'attributor context': it is a matter of what the interests and purposes are of those in the context in which the causal explanation is offered.

We can now appreciate how VK entails a version of contextualism. VK says that knowledge is true belief caused by intellectual virtue (i.e., 1). Because causal attributions are context-sensitive

(i.e., 2), and knowledge attributions involve causal attributions, knowledge attributions are context-sensitive (i.e., 3). Greco (2004, 2008) labels the resulting position virtue contextualism.

If virtue contextualism is correct, there will be conversations involving knowledge ascriptions that exhibit the contextual factors just described. Greco provides an illustrative case in which there is disagreement about whether someone ‘knows’ that is explained by the different emphases different conversational participants place on the importance of virtue in the production of a true belief:

CONCERT RIDE

Mother: How did you get home from the concert tonight, Honey?

Teenager: I got a ride home from a guy that I met there.

Mother: Are you crazy!? You didn’t know that he wasn’t a murderer!

Teenager: Don’t be so dramatic, Mom. We talked for, like, an hour. I could see he was really nice. Of course, I knew he wasn’t a murderer!

Mother: I can’t believe you! You’re lucky you didn’t end up on the side of the road somewhere.

Teenager: Mom! I think I’m a better judge of character than that. (Variation on Greco 2004: 386).

Assume that the driver is not a murderer, and so the teenager’s belief that he is not is true. The teenager’s judgement of the driver’s character is a necessary part of the explanation as to why she formed a true belief. But her true belief is not solely a result of her judgement of character. As Greco (2008: 419) puts it, ‘the world must co-operate’. If the driver were not in such a good mood, or if she’d waited just a second longer and asked the next driver instead, she might not have been so lucky, and her belief may have turned out false. Given that the teenager and mother

have different interests and purposes, they both emphasize different causal factors as being salient. The teenager is interested in getting home, and this means that she emphasizes the fact that she judged the driver to be really nice to be the salient causal factor that produced her true belief. The mother is interested in the safety of her teenager, and this means that she emphasizes the lucky cooperation of the world, and downplays virtue, as the salient causal factor that produced her daughter's true belief. Thus, the teenager and mother disagree about whether the teenager 'knows' because their different interests and purposes lead them to take different causal factors to be salient. Greco (2004: 386 - 387) take cases like *CONCERT RIDE* to be evidence for virtue contextualism.

Where does virtue contextualism fit in the taxonomy of different forms of epistemic contextualism given in chapter 1? Conversational contextualism says that the truth-conditions of knowledge claims depend upon the dynamics of the conversation in which the claim is made. Virtue contextualism says that the truth-value of 'S knows that p' varies depending upon the interests and purposes of those in the ascriber's context. What makes causal talk context-sensitive is the interests and purposes of speakers, not conversational dynamics. Thus, virtue contextualism is not a version of conversational contextualism. Stakes contextualism says that the truth-conditions of a knowledge claim vary with how much is practically at stake for those in the ascriber's context. What's practically at stake for a person is a function of their interests and purposes. Thus, we might think that virtue contextualism is a species of stakes contextualism. But virtue contextualism doesn't quite match up with the description of stakes contextualism given in chapter 1. There I said that stakes contextualism has it that the more that's at stake the more alternatives that need to be ruled out. Virtue contextualism says that the truth-values vary with interests and purposes because they determine that different causal factors are salient. Given this, perhaps it'd be best to amend the taxonomy by saying that there's a broad kind of

contextualism that we might call interests contextualism of which both stakes contextualism and virtue contextualism are different versions.

3. Virtue Contextualism Evaluated

The aim of this section is to evaluate how well Greco's virtue contextualism deals with the problem cases we've been considering. To anticipate, it comes up short. But before we can jump into this evaluation, it's important to notice that virtue contextualism is incomplete as it stands. It's unclear what verdict virtue contextualism reaches about the truth-value of knowledge claims in cases of disagreement and it's unclear how it arrives at a verdict. In **CONCERT RIDE** does the teenager speak truly? The mother? Both? Neither? It's important for virtue contextualism to have a story about how the truth-values of knowledge claims are determined in cases of disagreement if it's to have any chance of delivering the intuitive verdicts on the problem cases. Without such a story we simply can't tell what verdict virtue contextualism will reach in our problem cases. We need to be able to say, for example, that in **DISMISSED FROM CLASS** Rashaan speaks truly when he ascribes 'knowledge' to himself and that the Assistant Principal speaks falsely when he denies that Rashaan 'knows'. As it stands, virtue contextualism doesn't have the resources to deliver these verdicts. I'll consider three ways to complete the story and show that none of them are successful.

3.1. Both speak truly: the simple contextualist approach

One option for Greco is to say that both the mother and the teenager speak truly. When the mother speaks it's her interests and purposes that determine 'the cause' of the teenager's belief. The mother's interests and purposes determine that the teenager's true belief is caused by the

fact that the world happened to cooperate, and so the mother's knowledge denial is true. When the teenager speaks it's her interests and purposes that determine 'the cause' of her belief. The teenager's interests and purposes determine that her true belief is caused by virtue, and so the teenager's knowledge ascription is true. Both the mother and the teenager speak truly.

This way of taking virtue contextualism is a non-starter insofar as the problem cases are concerned. Getting the right verdict in these cases requires delivering the verdict that one person speaks truly and the other speaks falsely. But on the version of virtue contextualism that we're presently considering both speakers speak truly. By way of illustration, take **DISMISSED FROM CLASS** which was our principal 'no testimonial injustice' case from chapter 5. In the case Rashaan suffers from testimonial injustice. Intuitively Rashaan 'knows', the Assistant Principal denies that Rashaan 'knows' because he cannot eliminate a skeptical error possibility, and the reason the Assistant Principal raises the error possibility is because he harbours a racist prejudice.

Applying the modified version of virtue contextualism we're considering, when Rashaan claims to 'know' it is his interests and purposes that determine 'the cause' of his belief. His interests and purposes are in seeing justice done when he suffers from racism. This makes the salient causal factor his sound judgement that is informed by his history of experiences of racism. Rashaan's true belief is caused by this factor, and so his knowledge ascription is true. When the Assistant Principal denies that Rashaan 'knows' it is his interests and purposes that determine the salient causal factors. His interests and purposes are in protecting his staff. While this does not lead him to explicitly pick anything out as 'the cause' of Rashaan's belief, it does lead him to clearly indicate that Rashaan's belief is not caused by intellectual virtue. Led by his interests and purposes, the Assistant Principal raises an ineliminable skeptical error possibility and subsequently denies that Rashaan 'knows'. This amounts to the Assistant Principal denying that Rashaan's belief is caused by sound judgement. We can imagine the Assistant Principal reasoning as follows: 'If you have a true belief that's caused by virtue, then you should be able to rule out

the error possibility. But you can't rule out the error possibility. So, you're belief is not caused by virtue' Given that it's the Assistant Principal's interests and purposes that determine what the salient causal factors are, Rashaan does not hold a true belief that's caused by intellectual virtue and so the Assistant Principal's knowledge denial is true. So, on this way of taking virtue contextualism both Rashaan and the Assistant Principal speak truly. Given that getting the right verdict on **DISMISSED FROM CLASS** requires delivering the verdict that Rashaan speaks truly and the Assistant Principal speaks falsely we're left with an incorrect result.

3.2. Neither speaks truly: the supervenionist approach

A second option for Greco is to adopt a version of Keith DeRose's (2004, 2009) supervenionist strategy that was discussed in chapter 3. Recall that on DeRose's view, the personally indicated epistemic standards of every conversational participant influence the truth-conditions of knowledge claims. On this view, 'S knows that p' is true iff S meets the personally indicated epistemic standard of every conversational participant, false iff S fails to meet the personally indicated epistemic standard of every conversational participant, and truth-valueless iff some standards are met but others are not.

Applied to virtue contextualism, the idea would be that 'S knows that p' is true iff S has a true belief caused by virtue relative to the salient causal factor of every conversational participant. 'S knows that p' is false iff S does not have a true belief caused by virtue relative to the salient causal factor of every conversational participant. And, 'S knows that p' is truth-valueless iff S has a true belief caused by virtue relative to the salient causal factor of some conversational participant(s) and S does not have a true belief caused by virtue relative to the salient causal factor of other conversational participant(s).

Applying the view to **CONCERT RIDE**, ‘The teenager knows that the guy that gave her a lift home is not a murderer’ is truth-valueless. The knowledge ascription is true relative the salient causal factor emphasized by the teenager. It is false relative the salient causal factor emphasized by the mother. And so, at the level of conversation it is indeterminate whether the knowledge ascription is true or false: it is truth-valueless.

Unfortunately, this way of taking virtue contextualism does not help to deal with the problem cases either. Again, delivering the right verdict in these cases requires showing that one person speaks truly and the other speaks falsely. But on the supervaluationist version of virtue contextualism the target knowledge claims will all come out as truth-valueless. Taking **DISMISSED FROM CLASS** as an example once again, ‘Rashaan knows that Ms. H’s decision to send him out was guided by racial prejudice’ is truth-valueless. The knowledge ascription is true relative to the salient causal factor from Rashaan’s perspective. It is false relative to the salient causal factor emphasized from the Assistant Principal’s perspective, and so at the level of conversation it is indeterminate. Once again, the wrong result.

3.3. One speaks truly the other falsely: The metalinguistic approach

One way that we can say that only one or the other of the parties to a disagreement speak truly is to give a metalinguistic account of the disagreement and then provide some criteria for settling the disagreement and seeing which of the parties best meets the criteria.

According to a metalinguistic account, disagreement is not to be explained in terms of the propositions that are semantically expressed by the speakers’ utterances, but rather in terms of the propositions that are pragmatically conveyed. These are propositions about how the term ‘knows’ should be used. Esa Diaz-Leon (2017), building on David Plunkett and Tim Sundell (2013), has argued in favour of this as a contextualist explanation of disagreement about ‘knows’.

Applied to *CONCERT RIDE*, when the mother says, ‘You didn’t know that he wasn’t a murderer!’ she is expressing that the teenager doesn’t ‘know’ relative to what she deems should be taken to be the salient causal factor: the fact that the world co-operated. When the teenager says, ‘Of course I knew he wasn’t a murderer!’ she is expressing that she does ‘know’ relative to when she deems should be taken to be the salient causal factor: her good judgement of character. Thus, according to the metalinguistic story the mother and the teenager are disagreeing because they are expressing different propositions at the pragmatic level, they are conveying different information about how ‘knows’ should be used. These two propositions are incompatible. This explains the appearance of disagreement.

Now, to get a verdict that one speaks truly while the other speaks falsely we need some way to determine whose proposal for what we should take to be the salient causal factor wins out. One option is to appeal to the consequences of accepting one proposal over the other, and then go with whichever proposal has the better consequences. The issue this way of settling which proposal to accept is that even though it might sometimes be clear which proposal will have better consequences, it will often be difficult to know which will have the better consequences. For example, would it be better to accept Rashaan’s proposal for what we should take to be ‘the cause’ of his belief, or would it be better to reject Rashaan’s proposal as the Assistant Principal does? It’s clear that rejecting Rashaan’s proposal will have some bad consequences. It will lead to Rashaan suffering all of the harms of testimonial injustice described in chapter 5. But accepting Rashaan’s proposal could have bad consequences too. Depending on how the school deals with racism amongst staff, accepting Rashaan’s proposal could mean that Ms. H. loses her job, and perhaps even her career in teaching altogether. And while it’s right that Ms. H. deserves some kind of punishment for her racism, it’s not so obvious that she deserves to lose her career, especially if she is epistemically blameless for the acquisition and maintenance of her prejudice because of being brought up in a deeply racist epistemic environment (See Begby (2021) for

extensive discussion of the possibility of epistemically blameless prejudice). Assuming that these are accurate descriptions of the consequences of accepting the proposals made by each of Rashaan and the Assistant Principal, it's not entirely obvious which set of consequences are worse, and so it's not obvious which proposal for what to take 'the cause' to be we should go with. Thus, while modifying virtue contextualism along these metalinguistic lines might help in some cases, it won't help us in cases where there are serious difficulties in determining the consequences of different proposals and in trading them off against one another.

4. A Different Sort of Virtue Contextualism: Virtue Responsibilist Contextualism

We've seen that on three different ways of developing the view, Greco's virtue contextualism isn't up to the task of dealing with our problem cases. Greco derives virtue contextualism from VK: the thesis that knowledge is true belief produced by intellectual excellence. But VK is not the only resource that virtue epistemology has to offer that might be useful in building a version of contextualism that avoids the problems faced by extant versions of conversational contextualism identified earlier in this thesis. Recall the distinction between virtue reliabilism and virtue responsibilism outlined at the start of section 2. Reliabilists say that virtues are whatever faculties help us reliably attain truth. Responsibilists say that virtues are acquired states of character that help us to realise some epistemic good. Most importantly for our purposes, reliabilists deny that any particular motivation is necessary for virtue. According to reliabilism, intellectual virtues can be either natural, e.g., vision, or acquired, e.g., the skills that one gains when taking a logic course. Consequently, reliabilists deny that the possession of virtue requires an acquired intellectual motivation (Sosa 1991: 278). Responsibilists (e.g., Montmarquet 1993, Zagzebski 1996, Baehr 2013) say that virtue does require a particular kind of motivation: the virtuous agent has an

acquired motivation to care appropriately about epistemic goods such as truth and understanding and is motivated to act and think in particular ways in order to attain these epistemic goods.⁴⁵

Let's unpack the structure of the responsibilist account of intellectual virtues in a little more detail. Responsibilists broadly agree (e.g., Montmarquet 1993, Zagzebski 1996, Roberts & Wood 2007, and Baehr 2013) that intellectual virtues have a two-part motivational structure. Firstly, each particular virtue involves a disposition to act and think in certain ways. For example, the open-minded person considers and takes seriously a range of alternative ideas. A particular virtue involves a disposition to be motivated by a particular set of these ends associated with it. Charlie Crerar (2017: 755) calls these the *proximate ends* of that virtue. Secondly, in order to be virtuous, the agent's motivation for proximate ends must themselves be grounded in a more fundamental motivation for epistemic goods, such as truth, knowledge, or understanding. Crerar (2017: 755) calls these the *ultimate ends* of epistemic virtue. Thus, the agent that possesses the virtue of open-mindedness is motivated to take seriously a range of viewpoints, and this is because they possess a more fundamental motivation for epistemic goods. Put more precisely:

In order for S to possess intellectual virtue V, S must be (a) motivated to act and think in ways characteristic of V (proximate ends), for the reason that S is (b) motivated to attain epistemic goods (ultimate ends).

Why focus on the motivational component of virtue responsibilism? One thing that unifies all of the cases that we've looked at that prove problematic for conversational contextualism is that they involve an agent raising an alternative for non-virtuous reasons. Trolling cases involve agents who raise alternatives in order to derail the conversation. Our principal authority problem case involves someone wielding their authority in order to get off the hook for bad behaviour. Note that authority problem cases might involve motivations other than trying to get off the hook for

⁴⁵ Although, Greco allows for a 'weak motivation condition'. See Greco's (2004) 'How to preserve your virtue while losing your perspective'.

bad behaviour, but they will all involve abuses of authority used to insist on ineliminable skeptical hypotheses that undermine intuitively true knowledge ascriptions. It's hard to imagine any such case involving virtuous motivations. Our principal adaptive presupposition case involves someone raising an alternative because of their misogynistic attitudes. 'No testimonial injustice' cases involve an alternative being raised for reasons of prejudice. On Fricker's (2007) account the perpetration of testimonial injustice is explained by the fact that a hearer lacks what she calls the virtue of testimonial justice: the disposition to neutralize the effect of prejudice on one's credibility judgements.

The motivational component of virtue responsibilism together with the above diagnosis of the problem cases suggests a modified version of Lewisian relevant alternatives contextualism that replaces the various problematic rules of relevance that we've encountered earlier in the thesis, i.e. Lewis' rule of attention and Blome-Tillmann's rule of pragmatic presupposition, with the following rule:

Rule of intellectual virtue (RV): any alternative that is raised in an act of intellectual virtue is relevant in C.

This gives us a modified version of Lewisian relevant alternatives contextualism that retains its core features but replaces the rule of attention with the rule of intellectual virtue. Call the resulting view *virtue responsibilist contextualism* (VRC). In full, virtue responsibilist contextualism is the conjunction of (L), (RV), and the rest of Lewis's rules of relevance:

(L) 'S knows that p' is true in C iff S rules out every not-p alternative that is relevant in C and S's belief that p is properly based.

The addition of a basing requirement is a modification of Lewis's view that is proposed by Blome-Tillmann (2014) and which was mentioned in chapter 4. It ensures that 'S knows that p' does not come out as true in cases where S's belief is based on inappropriate epistemic grounds, such as

being based on palm readings or on tealeaves. Blome-Tillmann (2014: 31) convincingly argues that such cases constitute counterexamples to Lewis's view.

And Lewis's other rules of relevance:

Rule of belief (RB): any not-p alternative that S believes, or should believe, is relevant in C.

Rule of actuality (RACT): any not-p alternative that is actual is relevant in C.

Rule of resemblance (RRE): any not-p alternative that saliently resembles another relevant alternative is relevant in C.⁴⁶

It's important to notice that conforming to RV does not require you to be in full possession of intellectual virtue and act from it, rather it only requires that your utterance is an intellectually virtuous act. Aristotle drew a distinction between actions from virtue and virtuous acts. Actions from virtue are those where the agent possesses the virtue and performs an action that exhibits the excellences comprising the virtue. An agent performs a virtuous act when they act as the virtuous agent would act on a particular occasion.⁴⁷ That is, an agent performs an intellectually virtuous act when their particular act stems from intellectually virtuous motives, the action is what the intellectually virtuous would do in the circumstances, and they achieve some epistemic good as result. The agent does not need to have a stable motive in order to perform an intellectually virtuous act, as the agent who is in full possession of virtue does, rather they just need to exhibit intellectually virtuous motivation on the occasion of their action. RV says that when an alternative

⁴⁶ Lewis (1996: 558 - 559) also offers three further 'permissive rules': the rules of reliability, method, and conservation. These rules make no difference to the shape of the view that I'm proposing, and so I've left them out of our discussion. As Ichikawa (2017: 23) notes, it's not entirely clear why Lewis included these rules in the first place.

⁴⁷ There is disagreement about how precisely to spell out the idea of a virtuous act. This disagreement exists in the debate amongst virtue ethicists on how best to develop a virtue ethical account of right acts. Rosalind Hursthouse (1991) proposes that 'an act is right iff it is what the virtuous agent would do in the circumstances'. Michael Slote (1995) proposes that an act is right iff it exhibits a virtuous motive'. Christine Swanton (2001) proposes an account on which an act is right iff it is a virtuous act where '(1) an action is virtuous in respect V iff it hits the target of virtue V; (2) an action is right iff it is virtuous overall'.

is raised in an intellectually virtuous act it is relevant. This means that a speaker does not have to fully possess intellectual virtue in order to successfully raise an alternative to relevance, rather what's required is that they act as the intellectually virtuous agent would act on that occasion.

If raising an alternative is to be an act of intellectual virtue, what are the virtues that are exhibited when one raises an alternative? One virtue that immediately comes to mind is open-mindedness. Different virtue epistemologists provide different accounts of open-mindedness, as discussed in chapter 6 (see Riggs 2019 for an overview). I will focus on Zagzebski's account (1996). To be open-minded one must be motivated to consider a variety of viewpoints, ideas, and alternatives (these are, in Crerar's terms, the proximate ends of open-mindedness), which is in turn grounded in a more fundamental motivation for epistemic goods, such as truth (these are, in Crerar's terms, the ultimate ends). The fact that the intellectual actions that are performed when one exercises open-mindedness include considering alternatives is what makes it one of the virtues for the virtue responsibilist contextualist to focus on. Being virtuously open-minded requires hitting the mean when performing the intellectual actions that are characteristic of it. The intellectual virtues lie at the mean between the two vices of excess and deficiency. The agent who is deficient in open-mindedness is dogmatic, they consider and entertain too few alternatives. The agent who lies at the extreme of open-mindedness is naïve, they consider and entertain too many alternatives. They consider alternatives that they should ignore. The open-minded person hits the mean, considering and ignoring alternatives appropriately.

Another virtue that is often exhibited when people raise alternatives is one that has been overlooked in the virtue epistemology literature: that of intellectual rigour. Whilst I can't develop a thoroughgoing account of intellectual rigour here, it will be useful to provide a sketch to get a more complete picture of virtuous ways of raising alternatives. The intellectually rigorous agent is one who undertakes a thorough examination of the evidence, alternate viewpoints, and alternate possibilities (these are the proximate ends of intellectual rigour), which is grounded in

a more fundamental motivation for epistemic goods like truth and understanding (these are the ultimate ends). The actions of the rigorous agent will often be much like the actions of the open-minded agent, but they come apart. Consider the philosopher who is steadfast in his commitment to his view who considers a range of alternatives and objections to his view because he wants the best possible defence. He is rigorous, but given that he is steadfast, he is not open-minded. What's more, the vices of rigour are different from the vices of open-mindedness. The agent who is deficient in rigour is careless. He brings inquiry to a close too quickly. The agent who possesses rigour in excess is overcautious. He lets inquiry go on too long, failing to bring inquiry to a close when he should. The virtuously rigorous agent hits the mean, not closing inquiry too soon, and not letting it go on too long, but rather brings it to a close at just the right time.

To get a better sense of the virtue of rigour it's helpful to view it as the virtue that's on display in Mill's (1977 [1859]) marketplace of ideas that he sees as emerging from the protection of free speech. Mill argued that we should permit free speech because it allows for a marketplace of ideas to flourish that delivers various epistemic goods such as truth and understanding. Firstly, it enables us to come to attain truth. If free speech is censored, and that speech expresses truths, then we miss out on exchanging false for true beliefs. Secondly, it enables us to avoid dogmatism and develop greater understanding. It's beneficial for people to express falsehoods because we can then develop arguments against them and form arguments in favour of the truth. We come to understand the reasons for why beliefs are true and what's wrong with arguments that support false beliefs. Thus, on Mill's view the protection of free speech facilitates rigorousness, which in turn delivers a variety of epistemic goods. There are surely other virtues that might be on display when a person raises alternatives, intellectual curiosity is another that comes to mind. But a sketch of two of the most prominent is sufficient for illustrating the kinds of virtues that are on display when agents raise alternatives that are ruled in by RV.

Given that virtue responsibilist contextualism ties the truth-conditions of knowledge claims to conversational dynamics, albeit in a highly constrained way, it qualifies as a version of conversational contextualism.

Virtue responsibilist contextualism has a lot to recommend it. It can handle all the of the cases that have been shown to cause trouble for extant versions of conversational contextualism in earlier chapters. Let's work through each of these problem cases in turn, starting with trolling cases. Firstly, VACCINE. In this case it is intuitive that you 'know' that vaccines reduce the chance of polio. Anne te Vaks, the troll, enters the conversation and raises the skeptical possibility that the real explanation for changes in rates of polio are changes in diet and increased hygiene standards rather than vaccines. You cannot rule this possibility out. Now, chapter 3 is not explicit about Anne te Vaks's motivations, but there are plenty of clues to suggest that Anne te Vaks is not behaving in an intellectually virtuous manner. One plausible motive is that Anne te Vaks sets out to derail the conversation. After all, this is precisely what makes a troll a troll: 'their presence in conversations derail the goals of their interlocutors' (see chapter 3). Though considering alternatives is what the open-minded agent does (though plausibly this is one alternative too many, and so makes Anne te Vaks miss the mean of open-mindedness), her considering the alternative is not motivated by an epistemic good like truth or understanding. This means RV does not rule it in. Nor is it ruled in by any of the other rules of relevance, and so Anne te Vaks's raising of the alternative does not undermine the truth of your knowledge ascription. A second plausible motive is that Anne te Vaks is being sincere, but that it is the fact that she buys into (the discredited) anti-vaccination conspiracy theory that leads her to express skepticism about the efficacy of vaccination. Intellectual vice is a plausible explanation for buying into conspiracy theories. Quassim Cassam (2016), for example, argues that conspiracy theorists often exhibit the intellectual vice of gullibility. If vice is what is motivating Anne te Vaks to raise

the skeptical possibility then it is not ruled in by RV, and so Anne te Vaks's raising of the alternative does not undermine your knowledge ascription.

Let's turn to our second principal trolling case, **BRIDGE**. **BRIDGE** is structurally different to all of the other problem cases considered. In every other case a skeptical possibility is raised which in turn undermines an intuitively true knowledge ascription. In **BRIDGE**, by contrast, an anti-skeptical possibility is raised and undermines an intuitively true knowledge denial. In the case Shane and Julia come across a rather shoddy looking bridge. It is intuitive that Shane's utterance of 'You don't know that we won't fall off' to Julia is true, given that they can't rule out the possibility that they will not fall off. Alyssa, the troll, comes along and optimistically suggests that Shane and Julia do not need to rule out the possibility that they won't fall off in order to 'know' that they won't. Given DeRose's gap view, this renders the intuitively true knowledge denial gappy, i.e., truth-valueless. Can virtue responsibilist contextualism deal with **BRIDGE**? Yes, it can. According to virtue responsibilist contextualism, whether an alternative is relevant depends on whether it is ruled in by the rules of relevance. The possibility that we're interested in is the possibility that Shane and Julia will fall off the **BRIDGE**. Shane and Julia claim that it's relevant, Alyssa claims that it is not. Notice that RV is a prohibitive rule: it tells us when an alternative cannot be ignored. Alyssa claims that the alternative can be ignored as irrelevant, and so prohibitive rules cannot tell in favour of Alyssa's insistence. Assuming that the falling off possibility isn't ruled in by any other rules of relevance, the question of the alternative's relevance comes down to whether Shane and Julie are virtuously motivated in raising the falling off possibility. I think the answer is, plausibly, yes. Surely, they're concerned about their safety when it comes to crossing the shoddy looking bridge, and this concern means that they're concerned about the truth of the bridge's safety. Truth is an epistemic good, and acts of raising an alternative that are motivated by truth will count as relevant by RV. Given that Shane and Julia can't rule out

the falling off possibility, their knowledge denial comes out as true. RV gets the right result. Thus, virtue responsibilist contextualism adequately deals with trolling cases.

Let's turn our attention to authority problem cases. These were considered in chapter 4 as a problem for Blome-Tillmann's presuppositional epistemic contextualism. **SEXUAL HARASSMENT*** was my principal authority problem case. In the case Sarah receives an email from her boss that constitutes sexual harassment. In conversation, her boss raises the possibility that somebody hacked into his email account and sent a sexually harassing email to Sarah. Because Sarah is unable to rule out the possibility, her intuitively true attribution of 'knowledge' to herself is undermined. Virtue responsibilist contextualism deals with this case straightforwardly. In the case the boss knows the truth, and so cannot be raising the possibility in a bid to attain the epistemic good of truth. Rather, the boss's motivation in raising the alternative is to cover up his bad behaviour. Given his motivation, the boss's alternative is not rendered relevant by RV, and it does not undermine Sarah's intuitively true knowledge ascription. The correct result.

Next up, adaptive presupposition cases. This class of cases were also considered in chapter 4 as a problem for Blome-Tillmann's presuppositional epistemic contextualism. **SEXUAL HARASSMENT**** was my principal adaptive presupposition case. The case is the same as **SEXUAL HARASSMENT*** apart from the fact that in this case it's Sarah's colleague, Steve, who raises the skeptical possibility that someone hacked into Sarah's boss's email account and sent the sexually harassing email to Sarah. We know that Steve is a misogynist. Beliefs characteristically held by misogynists include beliefs that women lie about sexual harassment, see sexual harassment where there is none, and overreact in circumstances that are described as sexual harassment. These sets of false prejudicial beliefs plausibly explain why Steve raises the alternative. An act of raising an alternative because of prejudice against women is clearly not an act that is virtuously motivated. Thus, the alternative is not made relevant by RV, and does not

undermine Sarah's intuitively true knowledge ascription. Another correct result for virtue responsibilist contextualism.

The final class of problem cases are 'no testimonial injustice' cases which were discussed in chapter 5. These pose a problem for all of the versions of conversational contextualism that we've considered. My principal 'no testimonial injustice' case is **DISMISSED FROM CLASS**. This is the easiest case of them all for virtue responsibilist contextualism to handle. Recall that in the case Rashaan suffers from testimonial injustice. Intuitively Rashaan 'knows', the Assistant Principal denies that he does 'knows' because Rashaan cannot eliminate a skeptical error possibility, and the reason the Assistant Principal raises the error possibility is because he harbours a racist prejudice. It's the fact that the Assistant Principal raises the skeptical error possibility because of prejudice that makes it easy for virtue responsibilist contextualism to handle. Fricker (2007) claims that failure to neutralise the effect of prejudice on one's credibility judgement is to lack the virtue of testimonial justice. Others go further: Zagzebski (1996), Cassam (2015), and Kidd (2016) all argue that prejudice is an epistemic vice. This means that RV does not rule the error possibility in. Nor is it ruled in by any of the other rules of relevance. Given that the structure of no testimonial injustice cases are such that they always involve the raising of an alternative for reasons of prejudice, and that prejudice is an epistemic vice, it's clear that virtue responsibilist contextualism can always avoid no testimonial injustice verdicts.

Virtue responsibilist contextualism, then, can handle all of the problem cases that have been shown to trouble extant versions of conversational contextualism in earlier chapters.

5. Virtue Responsibilist Contextualism vs. Conversational Contextualism

Virtue responsibilist contextualism handles all of the cases that have been identified as problems for extant versions of conversational contextualism. This is good news for the view. But, for virtue

responsibilist contextualism to be preferable to conversational contextualism all-things-considered it needs to be shown that endorsing virtue responsibilist contextualism doesn't require forgoing the key benefits of extant versions of conversational contextualism. I'll look at two such benefits of conversational contextualism that have been discussed earlier in the thesis and show that they're shared by virtue responsibilist contextualism. Firstly, I'll show that virtue responsibilist contextualism can account for important data points that motivate conversational contextualism, such as DeRose's bank cases. Secondly, I'll show that virtue responsibilist contextualism, like Blome-Tillmann's (2014) presuppositional epistemic contextualism, has the resources to meet the objection that befell Lewis's (1996) view discussed in chapter 4, namely, that it makes satisfying 'knows' is too hard.

As we know by now, contextualism is frequently motivated by its ability to explain intuitions about pairs of cases like DeRose's bank cases:

LOW: Hannah and Sarah are discussing whether the bank is open on Saturday. Hannah recalls 'I've been in there on a Saturday before'. She says, 'I know that the bank is open on Saturday'. Sarah agrees.

HIGH: Hannah and Sarah are discussing whether the bank is open on Saturday. Hannah says, 'I've been in there on a Saturday before'. But Sarah replies 'Banks can change their opening hours. You don't know that the bank is open on Saturday'. Hannah agrees.

(variation on McKenna 2017: 321, variation on DeRose 1992).⁴⁸

Assuming that Hannah is in the same epistemic position in both cases, and that the bank is open, DeRose takes it that we will intuitively judge that both Hannah's knowledge ascription in LOW and Sarah's knowledge denial in HIGH are true. Contextualism is able to explain these judgements: the truth of 'S knows that p' requires that S rule out different sets of alternatives (or,

⁴⁸ These are the same versions of the cases given in chapters 1, 3, & 6.

meet different epistemic standards on DeRose's preferred description) in different contexts of ascription. The ability of extant versions conversational contextualism to explain judgements about LOW and HIGH rely upon the idea that speakers can change which alternatives are relevant (or, what epistemic standards are operative) by raising them to salience (e.g., DeRose 1992, 1995, 2004; Lewis 1996). When in HIGH Sarah says 'Banks do change their opening hours' she makes relevant a possibility that must be ruled out if it's to be truly said that Hannah 'knows' that the bank is open.

Can virtue responsibilist contextualism account for these intuitive judgements about the bank cases? A 'yes' answer to this question requires ensuring that the alternative in which the bank changes its opening hours is not relevant in LOW, but that it is relevant in HIGH.

Virtue responsibilist contextualism can deliver the verdict that Hannah 'knows' in LOW: there are no relevant alternatives that are ruled in by the rules of relevance endorsed by virtue responsibilist contextualism, that Hannah's evidence does not rule out. But can virtue responsibilist contextualism deliver the verdict that Hannah does not 'know' in HIGH? Is the alternative that the bank has changed its opening hours ruled in by any of the rules of relevance? It's not ruled in by belief, actuality, or resemblance. So, if it is to be ruled in it's RV that must do the work. The question then becomes, is Sarah's act of raising the possibility that the bank has changed its opening hours an act of intellectual virtue? The difficulty of providing a 'yes' answer here is that even if we can say that a virtuous agent would raise an alternative in the circumstance, we cannot tell whether Sarah is raising the alternative because she is motivated to attain an epistemic good: her behaviour is not a perfect guide to her motivations. This is the difficulty facing virtue responsibilist contextualism that I mentioned in the introduction to his chapter. But the fact that we do not have a perfect guide to Sarah's motivations does not mean that we have no indications whatsoever. One such indication is whether a person is acting in accordance with the Gricean (Grice 1989) rule of cooperativeness:

‘Make your conversational contribution as is required, at the stage at which it occurs, by the accepted purpose or direction of the talk exchange in which you’re engaged’ (Grice 1989: 27).

Grice makes use of this principle to develop a theoretical account of conversational implicature, but he thought that it is a quite general rule that governs all cooperative endeavours. Conformity with or departure from the rule provides an indication of whether or not a speaker has virtuous motivations. If an act of raising an alternative is in accordance with the rule of cooperativeness, then it’s a good indication that it’s an act of intellectual virtue. If an act of raising an alternative is not in accordance with the rule of cooperativeness, then it’s a good indication that it’s not an act of intellectual virtue.

What it takes for an act to be intellectually virtuous varies from context to context. One way in which acts of intellectual virtue are context-sensitive is that virtuous actions are appropriately sensitive to the purposes of the conversation. Raising brain-in-vat possibilities is virtuous (open-minded, rigorous, etc.) in an epistemology seminar on skepticism, but not in a courtroom. This means that if one’s epistemic actions are not appropriately sensitive to the goals of the conversation when raising an alternative, one is not exhibiting intellectual virtue. If one is being appropriately sensitive to the goals of the conversation when one raises an alternative, then this is an indication that one is exhibiting intellectual virtue. Applied to **HIGH**, the topic of discussion is whether the bank is open. Raising possibilities in which the bank has changed its opening hours is in line with the purposes of the conversation, and so we have an indication that Sarah’s act of raising an alternative is virtuous. Thus, the alternative is ruled in, and virtue responsibilist contextualism can deliver the verdict that the knowledge denial is true.

Let’s turn now to one of the central selling points of Michael Blome-Tillmann’s (2009, 2014) presuppositional epistemic contextualism: that it is an improvement on Lewis’s (1996) contextualism because it makes ‘knows’ less difficult to satisfy. Recall from chapter 4 that because

of the inclusion of the rule of attention which says that any alternative that is attended to is relevant, Lewis's view makes 'knows' too difficult to satisfy. This problem was identified and illustrated by Michael Williams (2001: 15) with the aid of the following case:

TEENAGER: Laurie is lying in bed. She hears her teenage son's bedroom window open, so she gets up and looks out. She sees someone who looks a lot like her son jumping out onto the ground below. She goes into the bedroom and finds it empty. She concludes that her son has slipped out, and spends the night fuming. In fact, this is exactly what happened. In the morning she challenges him: 'I know you went out late last night. I saw you'. The son responds: 'How do you know? It's possible that you dreamt the whole thing.'

On Lewis's view, the rule of attention means that the dreaming alternative is relevant, and so the son speaks truly in denying that Laurie 'knows'. But this is clearly the wrong result, of course Laurie 'knows'! And so, the rule of attention should be rejected. Blome-Tillmann's PEC delivers the right result on **TEENAGER** by replacing the rule of attention with rule of pragmatic presupposition which says that it's not enough to make an alternative relevant that attention is paid to it, but it must also be compatible with the pragmatic presuppositions of every conversational participant. Because the dreaming alternative is not compatible with the pragmatic presuppositions of Laurie and her son after Laurie makes clear that she will not take it seriously, it is not relevant and so does not undermine the truth of Laurie's knowledge claim. Thus, PEC delivers the right verdict on **TEENAGER**.

Can virtue responsibilist contextualism also deliver the right verdict on **TEENAGER** and avoid the issue of making 'knows' too hard to satisfy? Let's first examine how virtue responsibilist contextualism handles the specific case of **TEENAGER** and then examine whether it has the resources to provide a more general solution to the problem of 'knows' being too difficult to satisfy. Virtue responsibilist contextualism does deliver the right verdict on **TEENAGER**. The

dreaming alternative raised by the son is not made relevant by RV, or by any of the other rules of relevance. It's clear that the son does not have intellectually virtuous motivations in raising the dreaming hypothesis. Rather, he is motivated by a desire to get off the hook for sneaking out. Given that the dreaming alternative is not relevant, Laurie's knowledge ascription is true. Thus, virtue responsibilist contextualism delivers the right result on TEENAGER.

Let's now turn to the question of whether virtue responsibilist contextualism has the resources to provide a more general solution to the problem of 'knows' being too difficult to satisfy. The prospects look good. To see this, recall that the virtues associated with considering alternatives include open-mindedness and rigour. Both the open-minded and the rigorous agent consider a variety of alternatives. But one can possess open-mindedness or rigorousness in excess. Such a person considers too many alternatives, considering alternatives that they should ignore. For example, the excessively open-minded or overly rigorous person might frequently reach for far out skeptical alternatives: 'you could be a brain in a vat', 'you could be in the matrix', 'you could be dreaming'. These are precisely the kinds of alternatives that make 'knows' too difficult to satisfy. But this is not a problem for virtue responsibilist contextualism. In most contexts such alternatives will always be one alternative too many. The person who raises them exhibits open-mindedness or rigorousness in excess, and so cannot be said to be performing acts of intellectual virtue when raising them. These alternatives will not be ruled in by RV, and thus do not threaten to undermine intuitively true knowledge ascriptions.

The ultimate success of virtue responsibilist contextualism will not be decided merely by its ability to handle these two issues. We'll also want to see that it can handle further issues, such as the lottery paradox. That is work that will have to be done at a later date. For now, we will have to settle for demonstrating that virtue responsibilist contextualism is promising enough for that to be work that's worth carrying out.

6. A Further Objection

Virtue responsibilist contextualism deals well with the cases that were shown to cause trouble for extant versions of conversational contextualism in the earlier parts of this thesis. It also retains some of the important benefits of extant versions of conversational contextualism. To finish up this chapter, I want to consider one additional objection to virtue responsibilist contextualism.

The objection is as follows: whereas Lewis's contextualism makes 'knows' too hard to satisfy, it might be objected that virtue responsibilist contextualism makes 'knows' too easy to satisfy. RV says that any alternative that is raised in an act of intellectual virtue is relevant. This opens up an easy route to 'knowledge': don't bother acting in intellectually virtuous ways when raising alternatives. By not acting in intellectually virtuous ways in raising alternatives we limit the number of alternatives that a subject has to rule out and make it easy to satisfy 'knows'. My response is two-fold. Firstly, if this is a problem, then it's a problem for all versions of conversational contextualism. Applied to the versions of conversational contextualism that we've been considering, the idea would be that 'knows' becomes easy to satisfy by speakers abstaining from raising error possibilities. Secondly, we might say that there is a difference in how praiseworthy the subject would be for having a true belief across the cases. You're more epistemically praiseworthy for having a true belief and ruling out a wide set of alternatives than for having a true belief and ruling out a narrower set of alternatives. This difference in praiseworthiness across the case is to be explained in terms in differences in the manifestation of intellectual virtue across the cases.

7. Summary and Preview

Chapters 3, 4, & 5 showed that conversational contextualism faces a variety of problems because it does not place adequate constraints on the effects of conversational dynamics on the truth-

conditions of knowledge claims. Extant versions of conversational contextualism are overly permissive in allowing for the conversational moves of nefarious actors to influence the content of 'knows'. The challenge, then, was to find a way of appropriately constraining the effects of conversational dynamics on the truth-conditions of knowledge claims such that the intellectually vicious moves of nefarious actors do not undermine the truth of intuitively plausible knowledge claims. I've proposed virtue responsibilist contextualism as a way of appropriately constraining conversational dynamics. Virtue responsibilist contextualism retains the core features of Lewisian contextualism, but replaces Lewis's rule of attention with the rule of intellectual virtue which says that when acts of raising alternatives are intellectually virtuous the alternatives raised are relevant. Virtue responsibilist contextualism bars the actions of nefarious actors from making alternatives relevant, and this makes it well placed to deal with the problems shown to befall extant versions of conversational contextualism earlier on in the thesis. I've also shown that virtue responsibilist contextualism retains two of the substantial benefits of extant versions of conversational contextualism and that it does not fair any worse than extant versions of conversational contextualism with respect to a further objection. More work is needed to show that virtue responsibilist contextualism retains other important benefits of extant versions of conversational contextualism and that it doesn't generate any further substantial problems, and so I won't commit myself to it outright just yet. But enough has been said to show that it is a view worth developing.

Chapter 8: Making Relevant Alternatives Contextualism Moral

1. Introduction

In chapter 7 I attempted to deal with the cases that cause problems for existing versions of conversational contextualism, discussed in the earlier chapters, by developing a new version of conversational contextualism. This was virtue responsibilist contextualism, which places constraints on the effects of conversational dynamics on the truth-conditions of knowledge claims by prohibiting the conversational moves of nefarious actors from making alternatives relevant. In this chapter I'll take up another way of dealing with the problem cases by developing a version of contextualism that divorces the truth-conditions of knowledge claims from conversational dynamics altogether and instead ties them to some extra-conversational feature of context. An appropriate extra-conversational feature is suggested by the fact that there is a lot at stake across our problem cases. Our two principal trolling cases involve individuals facing high stakes decision scenarios. **VACCINE** involves an individual facing a decision of whether or not to act to vaccinate their child. **BRIDGE** involves two individuals facing the decision as to whether they should take the risk of crossing a rather dangerous looking bridge. Our principal authority problem, adaptive presupposition, and 'no testimonial injustice' cases all involve the subject of a knowledge ascription being exposed to epistemic wrongdoing as a result of the consideration of various skeptical possibilities. Numerous contextualists take stakes to influence the truth-conditions of knowledge claims (e.g., DeRose 1992; Cohen 1999; McKenna 2011). But the kind of stakes that are present in some of our problem cases are different in kind from the practical stakes that contextualists have traditionally recognised: they are moral rather than practical. This means that existing versions of stakes contextualism are not up to the task of handling our problem cases. Rather, what's required is a version of stakes contextualism that allows for moral, in addition to practical, stakes to influence the truth-conditions of knowledge claims. Taking inspiration from

moral encroachment theorists (e.g., Fritz 2017; Moss 2018; Schroeder 2018; Basu 2019; Bolinger 2020) who have argued that moral considerations can make a difference to whether a subject's belief is justified or whether they have knowledge itself, I'll propose a version of stakes contextualism that gives a role to moral stakes in determining the truth-conditions of knowledge claims and argue that it handles the problem cases. I'll call the view moral stakes contextualism. Just as with virtue responsibilist contextualism, I'll develop moral stakes contextualism in a Lewisian relevant alternatives framework. In the same way that stakes contextualists have said that high practical stakes can expand the set of relevant alternatives (or ratchet up standards), moral stakes contextualism says that high moral stakes can make more alternatives relevant. A further way in which moral stakes contextualism departs from existing stakes views is that it says that moral stakes not only exert pressure to make more alternatives relevant, but that they can also exert pressure to render alternatives irrelevant.

I proceed as follows: in section 2 I motivate the idea that the extra-conversational feature to focus on when developing a non-conversation based contextualism that deals with our problem cases is moral stakes. In section 3 I'll follow through on this motivation by proposing moral stakes contextualism and show that it successfully handles our problem cases. In section 4 I'll consider some difficult questions and objections for the moral stakes contextualist. This will help to work out some of the details of the view. I'll finish up in section 5 by carrying out a cursory comparison of the relative merits of virtue responsibilist contextualism and moral stakes contextualism.

2. From Salience to Stakes

In this section I'll motivate the idea that the extra-conversational feature to fix on when developing a non-conversation based contextualism that deals with our problem cases is moral

stakes, before going on to work out the details of the view that I'll call moral stakes contextualism in the next section.

One longstanding contextualist answer to the question of which features of context determine variations in the truth-conditions of knowledge claims is practical stakes. On this kind of view, if whether *p* is true is a matter of great practical significance, then epistemic standards are demanding; more demanding than if the truth of *p* is unimportant.⁴⁹ This is the view that I labelled stakes contextualism in chapter 1. To illustrate, consider how DeRose sees stakes as figuring in an explanation of intuitive judgements of his original formulations of famous bank cases:

LOW BANK: My wife and I are driving home on a Friday afternoon. We plan to stop at the bank on the way home to deposit our paychecks. But as we drive past the bank, we notice that the lines inside are very long, as they often are on Friday afternoons. Although we generally like to deposit our paychecks as soon as possible, it is not especially important in this case that they be deposited right away, so I suggest that we drive straight home and deposit our paychecks on Saturday morning. My wife says, 'Maybe the bank won't be open tomorrow. Lots of banks are closed on Saturdays.' I reply, 'No, I know it'll be open. I was just there two weeks ago on Saturday. It's open until noon.'

HIGH BANK: My wife and I drive past the bank on a Friday afternoon, as in **LOW BANK**, and notice the long lines. I again suggest that we deposit our paychecks on Saturday morning, explaining that I was at the bank on Saturday morning only two weeks ago and discovered that it was open until noon. But in this case, we have just written a very large and very important check. If our paychecks are not deposited into our checking account before Monday morning, the important check we wrote will bounce, leaving us in a very bad situation. And, of course, the bank is not open on Sunday. My wife reminds

⁴⁹ I've switched to the epistemic standards, as opposed to the relevant alternatives, model here because I am taking DeRose's discussion of stakes as my starting point.

me of these facts. She then says, ‘Banks do change their hours. Do you know that the bank will be open tomorrow?’ Remaining as confident as I was before that the bank will be open then, still, I reply, ‘Well, no, I don’t know. I’d better go in and check.’ (DeRose 1992: 913; 2009, 1 - 2).

The husband’s epistemic position with respect to the proposition that the bank will be open tomorrow is the same in both cases. Intuitively, DeRose claims, the husband’s knowledge ascription in LOW and his knowledge denial in HIGH are both true. What explains why both judgements are true is that different standards are operative in each of the contexts, where these different standards are the result of the different practical situation the husband finds himself in. Because it doesn’t matter much whether he’s right about the bank’s opening hours in LOW, the standards are low, and his knowledge ascription comes out as true. Because it does matter a great deal whether he’s right about the bank’s opening hours in HIGH, the standards are high, and his knowledge denial comes out as true. The higher the practical stakes, the higher the epistemic standards.⁵⁰

Plenty of other contextualists have also taken practical stakes to have a role in determining the truth-conditions of knowledge ascriptions. Stewart Cohen (1999: 58) sees stakes playing the same role as DeRose does⁵¹, and appeals to stakes in an explanation of intuitive judgements of his

⁵⁰ Though DeRose is often thought to hold the view that both stakes and conversational dynamics make a difference to the truth-conditions of knowledge claims, where the stakes element is as just described and the conversational element is as described in chapter 3, the relationship that DeRose sees between stakes and standards is not entirely clear. In his it seems clear that he does hold the view of the relationship between stakes and standards as just described. When commenting on what features of the context affects the truth-conditions of knowledge ascriptions, DeRose writes that he has ‘preferences for the more ‘objective’ features – like the importance of being right and what has been said in the conversation’ (1992: 916). In more recent work (2005: 177 - 178) he writes that, ‘I myself do not believe that such a difference in stakes is necessary for a difference in semantic standards for ‘knows’. In fact, I think that speakers are free to use standards even wildly inappropriate to the practical situation they face – for instance, to use low standards where they face an extremely high stakes situation in which it would be much wiser for them to employ much higher standards.’ Here DeRose seems to be suggesting that whilst stakes might make a particular epistemic standard appropriate, it’s ultimately down to speakers’ which standard they choose however (in)appropriate. This point of scholarship does not matter a great deal for the project of this chapter, but it is worth drawing attention to it so as to avoid misattributing a view to DeRose. What’s important to recognize for our purposes is that one feature of the context of utterance that has been said to make a difference to the truth-conditions of knowledge ascriptions is practical stakes.

⁵¹ The previous footnote notwithstanding.

famous airport cases. Robin McKenna's (2011) *interests contextualism* is perhaps the most thoroughgoing version of the stakes view in that it eschews giving a role to conversational dynamics altogether. Interests contextualism is a version of Lewisian relevant alternatives contextualism that retains the core elements of Lewis's view but replaces the rule of attention with what McKenna calls the 'Rule of relevance'. The rule of relevance says that an alternative is relevant if an ascriber has a reason to consider it (2011: 745). An ascriber has a reason to consider an alternative in which *p* iff their practical interests and purposes are such that they should consider the alternative (2011: 745). Given the husband's practical situation in **LOW BANK** – it doesn't matter very much whether he deposits his paycheck right away – he does not have a reason to consider the alternative that the bank has changed its opening hours, and so it is not relevant. Given the husband's practical situation in **HIGH** – it matters a great deal that he deposits his paycheck before Monday morning – he does have a reason to consider the alternative the bank has changed its opening hours, and so it is relevant. Interests contextualism is a stakes view because the more pressing your practical situation the more alternatives that you have a reason to consider. Though there are some important differences between these various versions of stakes contextualism, what's important for us to note is the role that they give to practical stakes is such that the more important it is to those in the context of ascription that a proposition be true, the more demanding the standards, or the more alternatives that need to be ruled out, in order for a knowledge ascription to be true. Hereafter I'll refer to this collection of views as 'practical stakes contextualism'.^{52,53}

⁵² Michael Blome-Tillmann (2009: 262 – 263) also takes stakes to influence the truth-conditions of knowledge ascriptions, albeit indirectly by affecting speaker pragmatic presuppositions. He says that whether one pragmatically presupposes *p* in a context will sometimes depend on stakes. If it is especially important whether the bank is open, one will not behave as if it is common ground that the bank has not changed its opening hours recently.

⁵³ David Lewis (1996: 556) expresses the temptation to say that stakes affect relevance when in a footnote he writes that he might well have endorsed a 'Rule of high stakes' according to which 'when error would be especially disastrous, few possibilities are properly ignored'. For Lewis, though, what matters is what's at stake for the subject of knowledge claims, rather than what's at stake for those in the context of attribution. Given that Lewis focuses on subject context stakes and that I'm interested in ascriber context stakes, I've left Lewis's view out of the main in text discussion.

Notice that there is a lot at stake across the problem cases that have been shown to cause trouble for conversational contextualism in the earlier parts of this thesis. In **VACCINE**, the proposition that vaccines reduce the chance of getting polio is high stakes in that your friend puts their child at risk and risks leaving their parental obligations unfulfilled by failing to act on it. In **BRIDGE**, the proposition that Shane and Julia won't fall off the bridge is high stakes in that they'll both be at great risk of injury, or even death, if they act on it. In both **SEXUAL HARASSMENT*** and **SEXUAL HARASSMENT****, the act of considering the proposition that somebody hacked Sarah's boss's email account exposes Sarah to epistemic wrongdoing as she is consequently denied 'knowledge' of her own experiences of sexual harassment, which in turn undermines her permissions to act on the proposition in the pursuit of justice. In **DISMISSED FROM CLASS**, the act of considering the proposition that Rashaan has misinterpreted Ms. H's actions exposes Rashaan to epistemic wrongdoing. He suffers testimonial injustice as he is wrongfully denied 'knowing' that he was a victim of racism.

So, contextualists recognise stakes as a factor that influences the truth-conditions of knowledge ascriptions, and all of the cases that have been shown to cause trouble for conversational contextualism are cases in which there is a lot at stake. This leads one to wonder whether stakes contextualism can handle the problem cases. Practical stakes contextualism handles **BRIDGE** straightforwardly. The practical stakes of acting on the proposition that Shane and Julia won't fall off the bridge are high for Shane and Julia. If they act on the proposition and it turns out to be false, then they'll both be at best severely injured and at worst will die. Given this, practical stakes contextualism can say that the epistemic standards are demanding; so demanding that Shane and Julia don't know that they won't fall off.⁵⁴ Unfortunately, **BRIDGE** is the only one of our problem cases that practical stakes contextualism can help us with. Whereas the stakes recognised by practical stakes contextualism are practical stakes, the stakes present in **VACCINE**, **SEXUAL**

⁵⁴ Or, on a relevant alternatives model, ineliminable alternatives in which Shane and Julia will fall off are relevant.

HARASSMENT*, SEXUAL HARASSMENT**, and DISMISSED FROM CLASS are moral. In VACCINE, what's at stake is the fulfilment of parental obligations relating to the health of the child. In SEXUAL HARASSMENT*, SEXUAL HARASSMENT**, and DISMISSED FROM CLASS, what's at stake is perpetration of epistemic wrongdoing against members of a marginalized group that has the downstream effect of removing permissions to act. If stakes contextualism is to account for our intuitive judgements about these cases, it must make space for moral stakes to influence the truth-conditions of knowledge claims.

The idea that moral stakes have a bearing on epistemology is starting to become more familiar due to the emergence of the moral encroachment literature. Moral encroachment is the thesis that the epistemic justification of a belief (and, according to some theorists, knowledge) can depend on moral considerations. The basic idea is that high moral stakes make it harder to be epistemically justified. COSMOS CLUB, originally used by Gendler (2011), is a central motivating case for moral encroachment:

COSMOS CLUB: The night before he is to be presented with the Presidential Medal of Freedom, John Hope Franklin hosts a celebratory dinner at the Cosmos Club, at which he is a member. All the other black men in the club are uniformed attendants. While walking through the club, a woman sees him, calls him over, presents her coat check and orders him to bring her coat. (Case as formulated in Bolinger forthcoming: 2).

The woman's belief is epistemically well supported: given the man is black and he is in this particular club, the probability that he is an attendant is extremely high. Yet, it seems that the woman is wrong to form the belief. Moral encroachers say that there are moral considerations that make the context a high stakes one, and that this means that the woman's evidence is not sufficient for it to be epistemically justified. Theorists disagree about exactly what makes the case high stakes. Basu (2019) says that given the exclusionary and subordinating history of the club, the woman would wrong Franklin to believe on the basis of his race that he was an attendant.

Schroeder (2018) says that the issue is one of forming a false belief that diminishes the status of its object. Moss (2018) says that the issue is causing harm with negative stereotyping inferences. Bolinger (2018) says that the issue is the harmful loss of opportunity for Franklin arising from the cumulative effects of many people making this inference. Exactly which feature of the case makes it high stakes is not important for our purposes. Rather, what matters is that the idea that moral stakes have a bearing on epistemic matters is well motivated. This fact, along with the presence of high moral stakes in our cases, and the fact that there is a long-standing recognition of the possibility of stakes influencing the truth-conditions of knowledge ascriptions, suggests that a parallel move to that which is made in the moral encroachment literature might also be made by contextualists to account for our problem cases.

The moral stakes that we want our version of stakes contextualism to account for are of two different kinds, each of which exerts a different kind of influence on the truth-conditions of knowledge claims. The first attaches to action. It says that there being high moral stakes of failing to act on a proposition can give us reason to ignore not-p alternatives. This is illustrated in **VACCINE**. In the case, the stakes of the parent not acting on the proposition that vaccines reduce the chances of polio puts the child at risk and leads to a failure to fulfil parental obligations. In order for stakes contextualism to be able to account for the intuitive verdict on **VACCINE** - that 'I know that vaccines reduce the chance of getting polio' is true - it needs to allow that high moral stakes can exert a pressure to disregard alternatives. This is a different kind of connection between stakes and action to that which is traditionally recognised by practical stakes contextualism which says that the higher the practical stakes of acting on a belief should it turn out to be false, the more demanding the standards for 'knows', or the more alternatives that are relevant.

The second kind of moral stakes is on display in **SEXUAL HARASSMENT***, **SEXUAL HARASSMENT****, and **DISMISSED FROM CLASS**. All three involve high moral stakes of

considering alternatives given that it exposes a subject of a knowledge claim to epistemic wrongdoing, and subsequently removes epistemic permissions to assert, act, and believe (assuming knowledge norms of assertion, action, and belief). Across all three cases an alternative is considered that the subject cannot eliminate, which in turn leads audiences to deny 'knowledge'. In *SEXUAL HARASSMENT** and *SEXUAL HARASSMENT*** Sarah is denied 'knowledge' of her experience of sexual harassment because she cannot eliminate the error possibility that somebody hacked Sarah's boss's email account and sent the sexually harassing email. This perpetuates the pattern in which sexual harassment victims are often disbelieved, and so reproduces injustice. In *DISMISSED FROM CLASS*, Rashaan is not believed because he cannot eliminate the error possibility that he is misinterpreting his teacher's actions when he judges her to be racist. This continues the pattern in which black people are judged to see racism where there is none, which also reproduces injustice. If a version of stakes contextualism is to deliver the correct verdicts on these cases, then it will need to make space for moral stakes to exert a pressure to disregard alternatives.⁵⁵

Let's take stock. Practical stakes contextualism says that that high practical stakes of acting on a proposition should it turn out to be false raises the standards, or expands the sphere of relevant alternatives, for 'knowledge'. The cases that have been shown to trouble conversational contextualism are cases in which there is a lot at stake, but in a different way to that which is recognised by practical stakes contextualism. Firstly, they contain moral stakes, not practical stakes. Secondly, the stakes are not merely associated with action, but also with the consideration of alternatives that expose subjects of knowledge claims to epistemic wrongdoing. Thirdly, whereas the practical stakes recognised by practical stakes contextualism exert a pressure to raise

⁵⁵ There are two recent exceptions to this failure to recognise the possibility of high stakes preventing standards from rising: Crewe & Ichikawa (forthcoming) note in passing that the moral costs of failing to believe someone's testimony can lower the degree of justification required for belief in a context. Against the backdrop of discussion about the routine discounting of the testimony of victims of sexual assault. This idea is noted only in passing, rather than being part of a fully developed theory. Basu (2019) tentatively suggests 'perhaps we have a moral burden to lower our evidential standards for believing the victim'. Basu is not working within a contextualist framework.

the standards or expand the set of relevant alternatives, the moral stakes present in the cases that we're focusing on exert a pressure to disregard alternatives. A stakes contextualism that deals with the problem cases we're concerned with will be one that can accommodate these various dimensions of stakes. In the remainder of this chapter I'll propose a view that does exactly that and defend it against a variety of objections.

3. Moral Stakes Contextualism

So, there are various ways in which the cases that cause trouble for conversational contextualism involve high moral stakes. Recognition of this fact provides motivation for a novel non-conversational contextualism that is well placed to deal with the problem cases. The basic thought that needs to be accommodated if there is to be a version of stakes contextualism that delivers the intuitive verdicts on our problem cases is that sometimes high moral stakes can raise the standards or make more alternatives relevant, and that other times high moral stakes can lower the standards or make certain alternatives irrelevant. Beyond this basic thought, it's a challenge to state exactly how moral stakes might influence standards or relevant alternatives. Writing in the context of the moral encroachment literature, Rima Basu has noted that:

‘I am skeptical, however, that we'll be able to articulate a general principle to determine in every case whether the threshold goes up or whether the threshold ought to go down.’

(Basu 2019: 13)

Nevertheless, I'll attempt one way of developing a moralised version of stakes contextualism. I'll call the view moral stakes contextualism. Moral stakes contextualism is a version of Lewisian relevant alternatives contextualism that drops Lewis's rule of attention, incorporates three of his original rules, and adds three further rules that tie the relevance of alternatives to different kinds of stakes. Moral stakes contextualism, then, is the conjunction of:

(L) 'S knows that p' is true in C iff S rules out every not-p alternative that is relevant in C and S's belief that p is properly based.

The relevant not-p alternatives are determined by the rules of relevance. Moral stakes contextualism includes three of Lewis's original rules of relevance:

Rule of belief (RB): any not-p alternative that S believes, or should believe, is relevant in C.

Rule of actuality (RACT): any not-p alternative that is actual is relevant in C.

Rule of resemblance (RRE): any not-p alternative that saliently resembles another relevant alternative is relevant in C.

And three further rules that ties relevant alternatives to stakes:

Rule of practical consideration (RPC): any alternative which someone in the context of ascription practically should consider is relevant in C (McKenna 2011: 745).

Rule of moral consideration (RMC): any alternative which someone in the context of ascription morally should consider is relevant in C.

Rule of moral disregard (RMD): any alternative which someone in the context of ascription morally should not consider is not relevant in C.⁵⁶

The rule of practical consideration is a modified version of McKenna's (2011) rule of relevance.

McKenna original version only takes into account the stakes of the ascriber. I've labelled the modified rule the rule of practical consideration in order to draw the contrast with the two moral

⁵⁶ The view that comes closest to my moral stakes contextualism is Alexander Guerrero's moral epistemic contextualism according to which 'how much one must do epistemically in evaluating whether to believe some proposition p varies depending on what actions or omissions a belief in p will justify.' Where the more morally significant the actions that a belief that p will justify, the more stringent the epistemic demands must be met before one is justified in acting as if p is true (2007: 68 - 69). Notice that Guerrero's view is concerned with how much epistemic labour one must undertake before one is *justified acting as if* p is true. He stops short of endorsing a full blown moral epistemic contextualism according which justified belief and knowledge are themselves subject to moral considerations.

rules. I've amended McKenna's original rule to take into account the practical stakes of everyone in the context of ascription. The inclusion of the rule of practical consideration means that moral stakes contextualism gives the same role to practical stakes as existing versions of practical stakes contextualism. This allows moral stakes contextualism to explain the intuitive verdicts on DeRose's two bank cases and on BRIDGE.

Let's unpack the rule of moral consideration. What does it mean to say that someone in the context of ascription morally should consider an alternative? Someone in the context of ascription morally should consider an alternative when failing to consider it brings with it a risk of perpetuating a moral bad. For example, when failing to consider not-p leads one to act on some false p which subsequently brings about a moral bad, or when failing to consider not-p leads one commit doxastic wrongdoing which is wronging someone in virtue of what we believe about them (see, e.g., Basu & Schroeder (2019) for discussion of the phenomenon of doxastic wronging).

Having non-epistemic reasons to engage in certain epistemic activities like considering alternatives is a normal feature of our normative lives. For example, Sarah Stroud (2006: 514) presents data that suggests that upon receipt of new information that a friend has done something objectionable, we should subject the evidence to greater scrutiny, extend more interpretive charity, and end up at different conclusions than if the information was about someone who is not a friend. One good explanation for this is that the demands of friendship generate non-epistemic reasons to engage in particular epistemic activities, like gathering more evidence and considering more explanations of the evidence before settling on a conclusion that casts a friend in a bad light, than if you were forming a belief about someone who's not a friend. Similar things can be said about how other roles and relationships that we find ourselves in generate non-epistemic reasons to engage in particular epistemic activities.

The rule of moral consideration is needed to account for the kinds of cases that are standardly taken to motivate non-contextualist versions of moral encroachment, like variants of COSMOS CLUB that involve knowledge claims. Consider a variation of COSMOS CLUB that involves an explicit knowledge claim. Suppose that it's true that John Hope Franklin is an attendant and the woman asserts, 'I'll give my coat to this man, I know he's an attendant'. Intuitively, the woman's knowledge ascription is false. What explains this? She has a true belief that's based on good evidence after all. The rule of moral consideration gives us an explanation of why the woman's knowledge ascription is false that parallels the moral encroachment explanation of why the woman's belief is not justified. Given the high moral stakes of forming a false belief, the alternative in which this man is an exception to the generalization is relevant. He is one of the few black members of the club. This is not an alternative that the woman's (statistical) evidence can rule out. So, by the rule of moral consideration, 'I know that he's an attendant' is false.

Let's now turn to the rule of moral disregard. What does it mean to say that someone in the context of ascription morally should not consider an alternative? Someone in the context of ascription morally should not consider an alternative when the moral stakes of considering the alternative are high. For example, when considering not-p leads one to refrain from acting on some true proposition which subsequently brings about a moral bad, or when considering not-p leads one commit epistemic wrongdoing.

The rule of moral disregard is the rule that allows us to account for the intuitive verdicts on VACCINE, SEXUAL HARASSMENT*, SEXUAL HARASSMENT** , and DISMISSED FROM CLASS. Firstly, VACCINE. The judgement that we want to vindicate is that the parent's utterance of 'I know that vaccines reduce the chance of getting polio' is true. Conversational contextualism, in particular, DeRose's gap view, fails to deliver this verdict because it says there is an ineliminable error possibility that is relevant, namely, the anti-vaxx possibility that's raised by Anne te Vaks. Moral stakes contextualism successfully delivers the verdict that the parent's

knowledge ascription is true because by the rule of moral disregard the anti-vaxx error possibility is irrelevant. The parent morally should not consider the possibility because doing is morally costly in that it gets in the way of the parent taking a course of action that has a high moral payoff. We can imagine the parent reasoning, 'well, if I can't rule out the anti-vaxx possibility, then I don't know that vaccines are efficacious, and if I don't know that vaccines are efficacious, then I shouldn't vaccinate my child'. Consideration of the anti-vaxx alternative leads to a failure to fulfil parental obligations and puts the child at risk. Because the anti-vax alternative is irrelevant, morally should not be considered, so not considering it does not undermine the parent's intuitively true knowledge claim. Thus, moral stakes contextualism delivers the right verdict on VACCINE.

We can take SEXUAL HARASSMENT*, SEXUAL HARASSMENT**, and DISMISSED FROM CLASS together because they all involve the same kind of moral stakes: considering an alternative exposes the subject of the knowledge claim to epistemic wrongdoing, which in turn undermines their ability to act in pursuit of justice. The verdicts that we want to deliver are that Sarah's and Rashaan's knowledge ascriptions are true (and that the knowledge denial uttered by Rashaan's Assistant Principal is false). Conversational contextualism fails to deliver these verdicts because it says that in each case there is an ineliminable error possibility that is relevant, namely, the possibility that someone hacked Sarah's boss's email account in SEXUAL HARASSMENT* and SEXUAL HARASSMENT** and the possibility that Rashaan misinterprets his teacher's actions in DISMISSED FROM CLASS. Moral stakes contextualism successfully delivers the verdict that all three knowledge ascriptions are true because, by the rule of moral disregard, together with the relevant moral claims, all of the undermining alternatives that are rendered relevant by conversational contextualism are rendered irrelevant by moral stakes contextualism. In SEXUAL HARASSMENT* Sarah's boss morally should not consider the possibility that somebody hacked his email account because by doing so he exposes Sarah to epistemic

wrongdoing (and undermines her ability to act in the pursuit of justice). In **SEXUAL HARASSMENT**** Steve morally should not consider the possibility that somebody hacked Sarah's boss's email account because by doing so he exposes her to epistemic wrongdoing (and undermines her ability to act in the pursuit of justice). In **DISMISSED FROM CLASS** the Assistant Principal morally should not consider the alternative in which Rashaan misinterprets his teacher's actions because doing so exposes Rashaan to epistemic wrongdoing. Because the alternatives morally should not be considered, they are irrelevant and do not undermine the intuitively true knowledge ascriptions. Thus, moral stakes contextualism delivers the right results in all of the morally problematic cases that have been adduced for conversational contextualism in the earlier chapters of this thesis.

4. Questions and Objections

Moral stakes contextualism handles all of the cases that have been identified as problematic for extant versions of conversational contextualism. This is good news for the view. But, for moral stakes contextualism to be a serious alternative to conversational contextualism it needs to be shown that it does generate any new substantive problems. In this section I'll consider a series of questions and objections to moral stakes contextualism in order to better work out the details of the view and test whether it faces any substantive problems.

Question 1. What happens when practical and moral considerations pull in different directions? Consider a variation of **SEXUAL HARASSMENT**** in which in addition to the high moral stakes that are present in the case, there are also very high practical stakes. Suppose that if Sarah acts on the proposition that her boss sexually harassed her, then her boss will lose his job, which is his only source of income, and his family will become homeless. Now we have a situation in which the moral stakes count in favour of disregarding the alternative in which someone hacked

Sarah's boss's email account, and in which the practical stakes count in favour of considering the alternative. By the rule of moral disregard, the alternative is irrelevant. By the rule of practical consideration, the alternative is relevant. Is the alternative irrelevant or relevant?

Answer. One option is to say that one or the other of the rules always takes priority. If morality always trumps, then by the rule of moral disregard the alternative is irrelevant, and Sarah 'knows'. If practical considerations always trump, then by the rule of practical consideration the alternative is relevant, and Sarah does not 'know'. This seems too blunt of a response. If morality trumps, then no matter how minor the moral considerations exerting pressure to disregard an alternative, and no matter how weighty the practical considerations exerting pressure to consider the same alternative, the alternative will be deemed irrelevant. This seems implausible. If practical considerations trump, then no matter how weighty the moral considerations exerting pressure to disregard an alternative, and no matter how minor the practical considerations exerting pressure to consider the same alternative, the alternative will always be deemed relevant. This too seems implausible. So much for the priority answer to the question.

A second option is to say that when the rule of moral disregard renders an alternative irrelevant and the rule of practical consideration renders an alternative relevant both rules operate just as they would if there was no conflict, with the former rule ruling the alternative out and the latter rule ruling the alternative in. The effect of this on the truth value of the knowledge ascription to which the alternative being discussed is both irrelevant and relevant, i.e., Sarah's utterance of 'I know that my boss sexually harassed me', is that it is rendered indeterminate. One cost of this option is that, assuming the knowledge norm of action, the indeterminacy of the knowledge ascription percolates out making it impermissible to act on the proposition that features in the knowledge ascription. Since if it's indeterminate whether you 'know', it's not true that you 'know', and so you cannot act and abide by the knowledge norm of action. This a recipe for inaction, which is a cost in particularly pressing practical situations.

A third option is to say that sometimes the weightiness of practical reasons will be such that the rule of practical consideration trumps, and that on other occasions the weightiness of moral reasons will be such that the rule of moral disregard trumps. On this picture, If the practical reasons to consider are stronger than the moral reasons to disregard, then the alternative is relevant. And if the moral reasons to disregard are stronger than the practical reasons to consider, then the alternative is irrelevant. I find myself most tempted by this option. It avoids some of the issues that come with the other two options: it respects the intuition that on some occasions practical considerations are more important than moral considerations and that on other occasions moral considerations are more important than practical considerations. While there will be indeterminacy in those cases where the weightiness of practical and moral considerations are perfectly equal, these will be far smaller in number than with option 2. This option does bring with it all of the difficulties of knowing how exactly to weigh reasons. But this is a very general difficulty that comes with trying to figure out how to weigh moral and practical reasons against each other. So, while option 3 is not entirely without difficulties it doesn't generate any new or special difficulties of its own.

Note also that this is not only a question for the moral stakes contextualist but also for proponents of pragmatic and moral encroachment. See Fritz (2017: 658 – 659) for discussion on the options for pragmatic and moral encroachers in responding to this question.

Question 2. What happens when moral considerations pull in different directions, where there are some moral reasons to consider and some moral reasons to disregard the same alternative? It's common to hear people say, 'there's a lot at stake on both sides'. Consider the debate over the proposition that transwomen are women. Gender critical feminists say that if we accept, and act on in certain contexts, this proposition and it turns out to be false, then we will have violated the rights of female-bodied people. So, by the rule of moral consideration, alternatives in which transwomen are not women should be taken seriously. Other feminists say that if we fail to accept,

and fail to act on in certain contexts, the proposition and it turns out to be true, then we will have perpetuated harm against transwomen. So, by the rule of moral disregard, alternatives in which transwomen are not women should be disregarded. So, it seems moral stakes contextualism both rules in and rules out the proposition.

Answer: I'm inclined to respond to this question in the same way that I responded to the previous question. Sometimes moral reasons on balance will tell in favour of disregarding an alternative, other times the moral reasons on balance will tell in favour of considering an alternative, and yet other times the moral reasons in favour of disregard and those in favour of consideration will be perfectly balanced such that it's indeterminate whether one 'knows'. Again, this is a question that faces all moral encroachers who hold that moral considerations can make the demands for knowledge, justification, or the satisfaction of 'knows' weaker or stronger, and so poses no special difficulty for moral stakes contextualism.

Objection 1. The wrong kinds of reasons objection. James Fritz (2019) has recently argued that some moral encroachment views, namely, the 'radical moral encroachment' views of Pace (2011) and Basu & Schroeder (2019) suffer from a wrong kind of reasons problem. These versions of moral encroachment say that the moral badness of forming a belief, such as when believing on inconclusive evidence that one's spouse has started drinking again undermines one's relationship (Basu & Schroeder 2019), counts against the epistemic rationality of forming the belief. Fritz argues that the moral badness of forming the belief is a reason that is of the wrong kind to count against the epistemic rationality of the forming the belief. One might think that there is a version of the wrong kind of reasons objection applied to moral stakes contextualism. The objection would be that only truth-related reasons, which are provided by evidence for or against the proposition in question, can affect whether one's epistemic position with respect to a proposition is good enough for 'knowledge that p'. As such, only truth-related reasons can affect whether an alternative is relevant. Moral reasons, being non-truth-related, cannot affect whether an

alternative is relevant. Moral reasons are the wrong kind of reasons to affect whether an alternative is relevant.

Response. This is a nonstandard way of levelling a ‘wrong kind of reasons’ objection. The standard way of levelling such an objection is to criticize an attitude for being based on the wrong kind of reasons or to criticize a theory for proposing that the wrong kind of reasons provide subjects with reasons for forming a particular attitude (see Gertken & Kieseewetter forthcoming for an overview of the structure of wrong kind of reasons objections). Fritz’s objection to radical moral encroachment takes this form. But moral stakes contextualism can’t be criticized in this way since it is no part of the view that a subject has moral reasons to form a belief. Rather, it says that moral reasons affect which alternatives a subject’s evidence must rule out if an ascription of knowledge to them is to be true. Beliefs themselves should be solely based on evidence. While it seems correct to say that belief shouldn’t be based on moral reasons, it’s less clear why which alternatives are relevant, should not be affected by moral considerations, such as the kinds of moral stakes that feature in our problem cases.

The nonstandard formulation also appears to carry little weight. As described, it would apply to every version of contextualism that gives a role to non-truth related considerations in the determination of the truth-conditions of knowledge ascriptions. Many versions of contextualism have this feature – all of those that give a role to practical stakes (DeRose 2009; McKenna 2011), those on which drawing attention to an alternative makes it relevant (Lewis 1996), those on which pragmatically presupposing an alternative makes it relevant (Blome-Tillmann 2009) – and have not been objected to on the basis of it, and so it can’t be that non-truth related reasons do not affect relevance.

Objection 2. The ignorance objection. This objection starts out by noting the very plausible assumption that we’re at least sometimes ignorant of what objective morality requires. This assumption entails that we’re also sometimes ignorant of whether or not an alternative is relevant,

which in turn entails that we're sometimes ignorant of whether 'knows' is satisfied. This implication, so the objection goes, is unacceptable.

Response. While we don't always know what objective morality requires, we often do. We do, for example, know that objective morality requires the fulfilment of parental obligations not to expose your children to undue risk, and that it requires not exposing victims of sexual harassment or racism to epistemic wrongdoing. Our ability to recognise these facts enables us see that they influence the relevance of alternatives in ways that enables moral stakes contextualism to get the right results on the cases that trouble conversational contextualism. This is a significant advantage of moral stakes contextualism.

It's also worth asking why the ignorance implication is unacceptable? It can't be because it's false that we're ever ignorant of whether 'knows' is satisfied. After all, we are sometimes ignorant whether 'knows' is satisfied. The following kind of exchange is a familiar one: Wife: 'Does James know that the bank is open on Saturday?' Husband: 'I don't know.' Without a compelling story about why ignorance in some set of cases is unacceptable the objection doesn't appear to have much force. Notice also that the purportedly problematic implication is an implication of all moral encroachment views, and it hasn't been taken to be decisive against those views.

Objection 3. The closing down inquiry objection. This is an objection to the rule of moral disregard. The objection says that by issuing moral reasons to disregard alternatives where considering them is high stakes, the rule of moral disregard prohibits inquiry into topics that should be permitted. Given that free inquiry is highly important, this implication gives us reason to be suspicious of the rule of moral disregard. Here's an example: the members of a science lab are discussing whether to amend its hiring policy to include an anti-discrimination component aimed at correcting for the fact that men outnumber women in STEM fields. The policy being enacted depends on the lab members knowing that discrimination is part of the explanation for men outnumbering women. The evidence possessed by the lab members doesn't strictly rule out

the possibility that men outnumber women in STEM fields because of biological differences. If the possibility is considered by the lab members, they will not be able to truly claim to ‘know’ that the disparity is due to discrimination and they won’t enact the anti-discrimination policy. This means that considering the possibility is high stakes: it undermines a course of action aimed at correcting for injustice. But suppose that some members of the lab have a research interest in investigating the possibility that differences between the number of men and women in STEM fields are to be explained in terms of biology. By the rule of moral disregard, the person can’t do this. This, so the objection goes, is an unacceptable closing down of inquiry.

If the closing down inquiry objection is compelling, then it can offer us an explanation as to why moral stakes can expand, but not contract, the range of relevant alternatives. Whereas expansion is consistent with further inquiry, which is a good thing, contraction prohibits further inquiry, which is a bad thing.

Response. A first point to note is that it’s simply not true that more inquiry is always better than bringing inquiry to a close. Sometimes the cost of further inquiry is just not worth the effort. Perhaps the issue at hand is just uninteresting. In **LOW BANK**, the husband could continue inquiring into whether the bank is open on Saturday, but what would be the point? Other times the pressing nature of our practical situation exerts pressure to bring inquiry to a close. In the time spent continuing to inquire into whether the bank is open tomorrow, it will have already closed today. A second point to note is that some philosophers (e.g., Pace 2011; Hannon 2020) have argued that practical and moral considerations are the only plausible means by which to settle whether inquiry has gone on long enough.

But it might be true that inquiry sometimes has to be brought an end by moral considerations, and yet still also be true that the rule of moral disregard objectionably shuts down inquiry in particular cases. One way to respond here is by embracing the implication. Perhaps it is just true that where further inquiry brings about a moral bad, we should not continue inquiring. As

valuable as free inquiry is, it is not always so valuable as to outweigh the value of avoiding the perpetration of a moral bad. We can supplement this response by noting that the cases that we're able to conjure up in making the closing down inquiry objection are ones in which it's unclear what the moral facts actually are and that this is what makes it seem as if closing down inquiry for moral reasons is objectionable. If this is right, then it can't be claimed that the rule of moral disregard does bring about an objectionable closing down of inquiry since whether or not this is true will only become clear once we've alleviated our moral ignorance.

This brings an end to the questions and objections. Moral stakes contextualism faces some difficult questions and objections, but enough has been said to show that none of them render the view a non-starter.

5. Virtue Responsibilist Contextualism or Moral Stakes Contextualism?

To finish up, I'll briefly compare the relative merits of virtue responsibilist contextualism and moral stakes contextualism.

Both views are well placed to handle the problem cases that have been shown to cause trouble for conversational contextualism. No surprise there, given that's what both views are designed to do!

Although I've argued that both views are able to account for important data points, such as the bank cases, that are key motivations for contextualism, moral stakes contextualism offers a less controversial account of the data. By including the rule of practical consideration, moral stakes contextualism accounts for the data in exactly the same way as traditional practical stakes versions of contextualism. By contrast, the virtue responsibilist explanation relies on drawing an inference from the behaviour of conversational participants to the presence of an unobservable motivation for epistemic goods. And while the behavioural evidence is good evidence, it is fallible evidence.

For this reason, moral stakes contextualism offers us an explanation of the data that is more secure.

On the other hand, virtue responsibilist contextualism perhaps has an edge insofar as it does not have to answer to some of the difficult questions and objections that moral stakes contextualism does. For example, virtue responsibilist contextualism does not appear to need to answer questions about conflicts between the rules of relevance that make up the view. Moreover, whereas moral stakes contextualism must answer to the objection that it shuts down inquiry by issuing moral requirements to disregard alternatives, virtue responsibilist contextualism does not and might even be thought to constrain inquiry in a desirable way by only allowing the contributions to inquiry of epistemically virtuous individuals to influence whether people 'know'.

This comparison is cursory at best. I think that what we learn from it is that both views have merits, and both have drawbacks, but that both are meritorious enough to deserve being more fully developed. But that is work that will have to wait for a later occasion.

Chapter 9: Conclusion

The project of this thesis has been to theorize about epistemic contextualism non-ideally. The strategy has been to remove the common conversational contextualist assumption that speakers have equal abilities to pull off conversational moves that affect the truth-conditions of knowledge ascriptions and replace it with the non-ideal assumption that speakers have unequal abilities to affect the truth-conditions of knowledge ascriptions, where these abilities vary with unfairly distributed power. Going non-ideal has turned up three payoffs. The first is that it has enabled us to identify unrecognized problems with existing forms of conversational contextualism. In chapter 3 we saw that Keith DeRose's gap view is vulnerable to trolling. I argued that this problem is recognizable even before going non-ideal, but that going non-ideal amplifies its severity. In chapter 4 we saw that Michael Blome-Tillmann's presuppositional epistemic contextualism is subject to two classes of counterexamples - authority problem cases and adaptive presupposition cases - and has a host of morally and politically concerning implications when combined with various plausible knowledge norms. In chapter 5 we saw that all the extant versions of conversational contextualism that we've considered are unable to vindicate our intuitions about clear cases of testimonial injustice. The problems adduced across chapters 3, 4, and 5 have a common core: they arise because the views criticized fail to place adequate constraints on how conversational dynamics influence the truth-conditions of knowledge ascriptions. The upshot is that an acceptable contextualism must either place adequate constraints on how conversational dynamics influence the truth-conditions of knowledge ascriptions, or divorce the truth-conditions of knowledge ascriptions from conversational dynamics altogether and instead tie them to some extra-conversational feature of context.

Chapter 6 delivered the second payoff of theorizing about contextualism non-ideally; it helped us to identify an unrecognized form of epistemic injustice, namely, irrelevance injustice.

Irrelevance injustice occurs either when a speaker raises an alternative that is not taken seriously when it should be, or when a speaker raises an alternative that is taken seriously when it should not be.

Chapters 7 and 8 delivered the third payoff of theorizing contextualism non-ideally. In those chapters I argued that the problems identified for extant versions of conversational contextualism in chapters 3, 4, & 5 motivate two new versions of contextualism. Motivated by recognition of the fact that all the problem cases adduced in chapters 3, 4, & 5 involve nefarious actors who raise alternatives for epistemically vicious reasons, in chapter 7 I proposed virtue responsibilist contextualism. Virtue contextualism is version of conversational contextualism that attempts to place adequate constraints on how conversational dynamics affect the truth-conditions of knowledge ascriptions. The main idea is that speakers can make alternatives relevant by raising them to salience, but only when they are virtuously motivated. Motivated by recognition of the fact that all the problem cases adduced in chapters 3, 4, & 5 involve high moral stakes, in chapter 8 I proposed moral stakes contextualism. Moral stakes contextualism divorces the truth-conditions of knowledge ascriptions from conversational dynamics entirely, and instead ties them to moral stakes. The main idea is that high moral stakes can sometimes make alternatives relevant and can sometimes make alternatives irrelevant. I have argued that both views are well placed to handle the problems that I have identified for extant versions of conversational contextualism, whilst retaining some of their key selling points. I have not argued for a full endorsement of either of these views. More work is needed to show that they do generate substantial problems of their own. But I hope to have done enough to show that this further work is worth undertaking.

Bibliography

Chapter 1 References

Blome-Tillmann, M. (2009). 'Knowledge and Presuppositions', *Mind*, 118 (470): 241 - 94.

Blome-Tillmann, M. (2014). *Knowledge and Presuppositions*, Oxford: Oxford University Press.

Briester, E. (2009). 'Feminist Epistemology, Contextualism, and Philosophical Skepticism', *Metaphilosophy*, 40 (5): 671 - 688.

Briester, E. (2017). 'Feminism and Contextualism', in *The Routledge Handbook to Epistemic Contextualism*, ed. J.J. Ichikawa. Taylor & Francis.

Davies, A. (2016). 'How to Silence Content with Porn, Context and Loaded Questions', *European Journal of Philosophy*, 24 (2): 498 - 522.

DeRose, K. (1992). 'Contextualism and Knowledge Ascriptions', *Philosophy and Phenomenological Research*, 52: 931 - 29.

DeRose, K. (1995). 'Solving the Skeptical Problem', *Philosophical Review*, 104 (1): 1 - 52.

DeRose, K. (2009). *The Case for Contextualism*, Oxford: Oxford University Press.

Dotson, K. (2011). 'Tracking Epistemic Violence, Tracking Practices of Silencing', *Hypatia*, 26 (2): 236 - 257.

Dretske, F. (1970). 'Epistemic Operators', *Journal of Philosophy*, 67 (24): 1007 - 23.

Fricker, M. (2007). *Epistemic Injustice: Power and the Ethics of Knowing*, Oxford: Oxford University Press.

Goldman, A. (1976). 'Discrimination and Perceptual Knowledge', *Journal of Philosophy*, 73: 771 - 91.

Goldman, A. (2007). 'Philosophical Intuitions: Their Target, Their Source, and Their Epistemic Status', *Grazer Philosophische Studien*, 74: 1 - 26.

- Hawthorne, J. & Stanley, J. (2008). 'Knowledge and Action', *Journal of Philosophy*, 105 (10): 571 – 590.
- Hesni, S. (2018). 'Illocutionary Frustration', *Mind*, 127 (508): 947 – 976.
- Ichikawa, J.J. (2020). 'Contextual Injustice', *Kennedy Institute of Ethics Journal*, 30 (1): 1 – 30.
- Ichikawa, J.J. (2017a). *Contextualising Knowledge: Epistemology and Semantics*, Oxford: Oxford University Press.
- Ichikawa, J.J. (2017b). 'Introduction: What is Epistemic Contextualism?' in *The Routledge Handbook to Epistemic Contextualism*, ed. J.J.Ichikawa. Taylor & Francis.
- Jones, K. (2002). 'The Politics of Credibility', in *A Mind of One's Own: Feminist Essays on Reason and Objectivity*, ed. Louise Antony & Charlotte Witt. Westview Press.
- Kelp, C. (2011). 'What's the Point of 'Knowledge' Anyway?', *Episteme*, 8 (1): 53 – 66.
- Langton, R. (1993). 'Speech Acts and Unspeakable Acts', *Philosophy and Public Affairs*, 22 (4): 293 – 330.
- Langton, R. & Hornsby, J. (1998). 'Free Speech & Illocution', *Legal Theory*, 4 (1): 21 – 37.
- Langton, R. & West, C. (1999). 'Scorekeeping in a Pornographic Language Game', *Australasian Journal of Philosophy*, 77 (3): 303 – 319.
- Langton, R. (2018). Blocking as Counter-Speech. In D.Fogal, D.W.Harris, & M.Moss (Ed.), *New Work on Speech Acts*. (pp. 144 - 162). Oxford: Oxford University Press.
- Lewis, D. (1979). 'Scorekeeping in a Language Game', *Journal of Philosophical Logic*, 8 (1): 339 – 359.
- Lewis, D. (1996). 'Elusive Knowledge', *Australasian Journal of Philosophy*, 74: 549 – 567.
- Maitra, I. (2009). 'Silencing Speech', *Canadian Journal of Philosophy*, 39 (2): 309 – 338.
- Maitra, I. (2012). 'Subordinating Speech', in *Speech and Harm: Controversies over Free Speech*, ed. M.K.McGowan & I.Maitra. Oxford: Oxford University Press.

- McGowan, M.K. (2009). 'Oppressive Speech', *Australasian Journal of Philosophy*, 87 (3): 389 – 407.
- McKenna, R. (2011). 'Interests Contextualism', *Philosophia*, 39: 741 – 750.
- McKenna, R. (2017). 'Conversational Kinematics', in *The Routledge Handbook to Epistemic Contextualism*, ed. J.J.Ichikawa. Taylor & Francis.
- Medina, J. (2013). *The Epistemology of Resistance*, Oxford: Oxford University Press.
- Pohlhaus, G. (2014). 'Discerning the Primary Harm in cases of Testimonial Injustice', *Social Epistemology*, 28 (2): 99 – 114.
- Schaffer, J. (2005). 'What Shifts? Thresholds, standards, or alternatives?' *Contextualism in Philosophy: Knowledge, Meaning, and Truth*, ed. G.Preyer & G.Peter. Oxford: Oxford University Press.
- West, C. (2012). 'Of Expression and Racist Hate Speech', in *Speech and Harm: Controversies over Free Speech*, ed. M.K.McGowan & I.Maitra. Oxford: Oxford University Press.

Chapter 2 References

- Blome-Tillmann, M. (2013). 'Contextualism and the Knowledge Norms', *Pacific Philosophical Quarterly*, 94 (1): 89 – 100.
- Brown, J. (2008a). 'Knowledge and Practical Reason', *Philosophy Compass*, 3 (6): 1135 – 1152.
- Brown, J. (2008b). 'Subject-sensitive Invariantism and the Knowledge Norm for Practical Reasoning', *Nous*, 42: 167 – 189.
- Cohen, S. (2004). 'Knowledge, Assertion, and Practical Reasoning', *Philosophical Issues*, 14 (1): 482 – 491.
- Craig, E. (1990). *Knowledge and the State of Nature*, Oxford: Oxford University Press.
- DeRose, K. (2002). 'Assertion, Knowledge, and Context', *Philosophical Review*, 111 (2): 167 – 203.

- DeRose, K. (2009). *The Case for Contextualism*, Oxford: Oxford University Press.
- Fantl, J. & McGrath, M. (2009). *Knowledge in an Uncertain World*, Oxford: Oxford University Press.
- Fricker, M. (2007). *Epistemic Injustice: Power and the Ethics of Knowing*, Oxford: Oxford University Press.
- Fricker, M. (2008). 'Scepticism and the Genealogy of Knowledge', *Philosophical Papers*, 37: 27 - 50.
- Gerken, M. (2017). *On Folk Epistemology: How do we Think and Talk About Knowledge*, Oxford: Oxford University Press.
- Gerken, M. (2019). 'Pragmatic Encroachment and the Challenge from Epistemic Injustice', *Philosophers Imprint*, 19 (15): 1 - 19.
- Hannon, M. (2015). 'The Importance of Knowledge Ascriptions', *Philosophy Compass*, 10 (12): 856 - 866.
- Hawthorne, J. (2004). *Knowledge and Lotteries*, Oxford: Oxford University Press.
- Hawthorne, J. & Stanley, J. (2008). 'Knowledge and Action', *The Journal of Philosophy*, 105 (10): 571 - 590.
- Ichikawa, J. (2012). 'Knowledge Norms and Acting Well', *Thought*, 1 (1): 49 - 55.
- Kim, B. (2017). 'Pragmatic Encroachment in Epistemology', *Philosophy Compass*, 12 (5): 14215
- Kusch, M. (2009). 'Testimony and the Value of Knowledge', in *Epistemic Value*, Eds. Haddock, A., Millar, A., & Pritchard, D. Oxford: Oxford University Press.
- Kappel, K. (2010). 'On Saying that Someone Knows: Themes from Craig', in *Social Epistemology*, Eds. Pritchard, D., Millar, A. & Haddock, A. Oxford: Oxford University Press.
- Kelp, C. (2011). 'What's the Point of 'Knowledge' Anyway?', *Episteme*, 8 (1): 53 - 66.

McGlynn, A. (forthcoming). 'Objects or Others? Epistemic Agency and the Primary Harm of Testimonial Injustice', *Ethical Theory and Moral Practice*.

Pohlhaus, G Jr. (2014). 'Discerning the Primary Harm in Cases of Testimonial Injustice', *Social Epistemology*, 28 (2): 99 - 114.

Pritchard, D. (2009). *Knowledge*, Basingstoke: Palgrave MacMillan.

Rysiew, P. (2012). 'Epistemic Scorekeeping', in *Knowledge Ascriptions*, Eds. Brown, J. & Gerken, M. Oxford: Oxford University Press.

Srinivasan, A. (2016). 'Philosophy and Ideology', *Theoria*, 31 (3): 371 - 380.

Stanley, J. (2005). *Knowledge and Practical Interests*, Oxford: Oxford University Press.

Unger, P. (1975). *Ignorance: The Case for Skepticism*, Oxford: Clarendon Press.

Williamson, T. (2000). *Knowledge and its Limits*, Oxford: Oxford University Press.

Chapter 3 References

Blome-Tillmann, M. (2013). 'Contextualism and the Knowledge Norms', *Pacific Philosophical Quarterly*, 94 (1), 89-100.

Brown, J. (2006). 'Contextualism and Warranted Assertability Manoeuvres', *Philosophical Studies*, 130: 407-435.

Buckwalter, Wesley (2017). 'Epistemic Contextualism and Linguistic Behavior', In J.J. Ichikawa (Ed.), *The Routledge Handbook to Epistemic Contextualism* (pp. 44-56). London: Routledge.

Cohen, S. (1988). 'How to Be a Fallibilist', *Philosophical Perspectives*, 13: 91-123.

DeRose, K. (1992). 'Contextualism and Knowledge Attributions', *Philosophy and Phenomenological Research*, 52, 913-29.

DeRose, K. (1995). 'Solving the Skeptical Problem', *The Philosophical Review*, 104 (1): 1-52.

DeRose, K. (2004). 'Single Scoreboard Semantics', *Philosophical Studies*, 119 (1 - 2): 1-21.

DeRose, K. (2009). *The Case for Contextualism*, Oxford: Oxford University Press.

Keefe, R. (2008). 'Vagueness: Supervaluationism', *Philosophy Compass*, 3 (2), 315-324.

Langton, R. (2018). 'Blocking as Counter-Speech', in D.Fogal, D.W.Harris, & M.Moss (Ed.), *New Work on Speech Acts*, (pp. 144 - 162), Oxford: Oxford University Press.

Lewis, D. (1996). 'Elusive Knowledge', *Australasian Journal of Philosophy*, 74 (4): 549-567.

McKenna, R. (2017). 'Conversational Kinematics', In J.J.Ichikawa (Ed.), *The Routledge Handbook to Epistemic Contextualism* (pp. 321-331). London: Routledge.

Richard, M. (2004). 'Contextualism and Relativism', *Philosophical Studies*, 119 (1-2), 215-242.

Williamson, T. (2000). *Knowledge and Its Limits*, Oxford: Oxford University Press.

Chapter 4 References

Abramson, K. (2004). 'Turning up the Lights on Gaslighting', *Philosophical Perspectives*, 28 (1): 1 - 30.

Beebe, J. (2012). 'Social Functions of Knowledge Ascriptions', in *Knowledge Ascriptions*, ed. Brown, J. & Gerken, M. Oxford: Oxford University Press.

Blome-Tillmann, M. (2009). 'Knowledge and Presuppositions', *Mind*, 118 (470): 241 - 94.

Blome-Tillmann, M. (2013). 'Contextualism and the Knowledge Norms', *Pacific Philosophical Quarterly*, 94 (1): 89 - 100.

Blome-Tillmann, M. (2014). *Knowledge and Presuppositions*, Oxford: Oxford University Press.

Blome-Tillmann, M. (2015). 'Ignorance, Presuppositions, and the Simple View', *Mind*, 124 (496): 1221 - 1230.

- Bolinger, R.J. (2020). 'Varieties of Moral Encroachment', in *Philosophical Perspectives*, 34 (1): 5 - 26.
- Craig, E. (1990). *Knowledge in the State of Nature*, Oxford: Oxford University Press.
- Crewe, B. & Ichikawa, J.J. (forthcoming). 'Rape Culture and Epistemology', in *Applied Epistemology*, ed. Lackey, J. Oxford: Oxford University Press.
- DeRose, K. (2004). 'Single Scoreboard Semantics', *Philosophical Studies*, 119 (1 - 2): 1 -21.
- DeRose, K. (2009). *The Case for Contextualism*, Oxford: Oxford University Press.
- Dormandy, K. (2018). 'Epistemic Authority or Proper Basing?', *Erkenntnis*, 83 (4): 773 - 791.
- Fritz, J. (2017). 'Pragmatic Encroachment and Moral Encroachment', *Pacific Philosophical Quarterly*, 98 (s1): 643 - 661.
- Hawthorne, J. & Stanley, J. (2008). 'Knowledge and Action', *Journal of Philosophy*, 105 (10): 571 - 590.
- Ichikawa, J.J. (2015). 'Ignorance and Presuppositions', *Mind*, 124 (496).
- Kelp, C. (2011). 'What's the Point of 'Knowledge' Anyway?', *Episteme*, 8 (1): 53 - 66.
- Langton, R. (2018). 'Blocking as Counter-Speech', in *New Work on Speech Acts*, ed. Fogal, D., Harris, D.W., and Moss, M. Oxford: Oxford University Press.
- Lewis, D. (1996). 'Elusive Knowledge', *Australasian Journal of Philosophy*, 74 (4): 549 - 67.
- Manne, K. (2017). *Down Girl: The Logic of Misogyny*, Oxford University Press.
- McKenna, R. (2017). 'Conversational Kinematics', in *Routledge Handbook of Epistemic Contextualism*, ed. J.J. Ichikawa. Taylor & Francis.
- Podosky, P.M.C. (forthcoming). 'Gaslighting, First and Second Order', *Hypatia*.
- Stalnaker, R. (1974). 'Pragmatic Presuppositions', in Milton Munitz & Peter Unger (eds.), *Semantics and Philosophy*. New York: New York University Press: 197 - 214.

- Stalnaker, R. (2002). 'Common Ground', *Linguistics and Philosophy*, 25 (5 - 6).
- Stark, C. (2019). 'Gaslighting, Misogyny, and Psychological Oppression', *The Monist*, 102 (2): 221 - 235.
- Williams, M. (2001). 'Contextualism, Externalism, and Epistemic Standards', *Philosophical Studies*, 103 (1): 1 -23.
- Williamson, T. (2000). *Knowledge and It's Limits*, Oxford: Oxford University Press.
- Zagzebski, L. (2013). 'A Defense of Epistemic Authority', *Res Philosophica*, 90 (2): 293 - 306.

Chapter 5 References

- Alcoff, L. (2005) *Visible Identities: Race, Gender, and the Self*, Oxford: Oxford University Press.
- Anderson, D.E. (2017). 'Conceptual Competence Injustice', *Social Epistemology*, 31 (2): 210 - 223.
- Barke, A. (2004). 'Epistemic Contextualism', *Erkenntnis*, 61 (2 - 3): 353 - 373.
- Blome-Tillmann, M. (2014). *Knowledge and Presuppositions*, Oxford: Oxford University Press.
- Brendel, E. (2005). 'Why Contextualists Cannot Know They Are Right: Self-Refuting Implications of Contextualism', *Acta Analytica*, 20 (2): 38 - 55.
- Brown, J. (2006). 'Contextualism and Warranted Assertability Manoeuvres', *Philosophical Studies*, 130: 407 - 435.
- Cohen, S. (1999). 'Contextualism, Skepticism, and the Structure of Reasons', *Philosophical Perspectives*, 13 (Epistemology): 57 - 89.
- DeRose, K. (1992). 'Contextualism and Knowledge Attributions', *Philosophy and Phenomenological Research*, 52: 913 - 29. 16
- DeRose, K. (2002). 'Assertion, Knowledge, and Context', *The Philosophical Review*, 111 (2): 167 - 203.

- DeRose, K. (2009). *The Case for Contextualism: Knowledge, Skepticism and Context, Vol. 1*, Oxford: Oxford University Press.
- DeRose, K. (2017). *The Case for Contextualism: Knowledge, Skepticism, and Context, Vol. 2*, Oxford: Oxford University Press.
- Dotson, K. (2012). 'A Cautionary Tale: On Limiting Epistemic Oppression', *Frontiers: A Journal of Women Studies*, 33 (1): 24 - 47.
- Dormandy, K. (2018). 'Epistemic Authority or Proper Basing?', *Erkenntnis*, 83 (4): 773 - 791.
- Fricker, M. (2007). *Epistemic Injustice: Power and the Ethics of Knowing*, Oxford: Oxford University Press.
- Hinchman, T. (2005). 'Telling as Inviting to Trust', *Philosophy and Phenomenological Review*, 70 (3): 562 - 587.
- Ichikawa, J.J. (2017). *Contextualising Knowledge: Epistemology and Semantics*, Oxford: Oxford University Press.
- Kornblith, H. (2000). 'The Contextualist Evasion of Epistemology', *Philosophical Issues*, 10: 24 - 32.
- Koksvik, O. (Manuscript). 'Benefitting from Epistemic Injustice'.
- Lewis, D. (1996). 'Elusive Knowledge', *Australasian Journal of Philosophy*, 74: 549 - 567.
- Lewis, D. (1979). 'Scorekeeping in a Language Game', *Journal of Philosophical Logic*, 8/1: 339 - 59.
- Medina, J. (2013). *The Epistemology of Resistance*, New York: Oxford University Press.
- Mills, C. (2007). 'White Ignorance', *Race and Epistemologies of Ignorance*, ed. S. Sullivan & N. Tuana. Albany, NY: SUNY Press. pp. 11 - 38.
- Montgomery, B. (2017). 'Epistemic Contextualism and the Shifting the Question Objection', *The Routledge Handbook of Epistemic Contextualism*, ed. J.J. Ichikawa. London & New York: Routledge Taylor & Francis Group.

- Peet, A. (2017). 'Epistemic Injustice in Utterance Interpretation', *Synthese*, 194 (9): 3421 - 3443.
- Pohlhaus, G. (2012). 'Relational Knowing and Epistemic Injustice: Towards a Theory of Willful Hermeneutical Ignorance', *Hypatia*, 27 (4): 715 - 735.
- Pynn, G. (forthcoming). 'Testimonial Injustice and Epistemic Degradation', in *Applied Epistemology*, ed. J. Lackey.
- Pynn, G. (2015). 'Pragmatic Contextualism', *Metaphilosophy*, 46(1): 26 - 51.
- Schiffer, S. (1996). 'Contextualist Solutions to Skepticism', *Proceedings of the Aristotelian Society*, 96: 317 - 333.
- Sosa, E. (2000). 'Contextualism and Skepticism', *Nous*, (s1) 1 - 18.
- Stalnaker, R. (1974). 'Pragmatic Presuppositions', in *Semantics and Philosophy*, ed. M. Munitz & P. Unger. New York: New York University Press.
- Stalnaker, R. (2002). 'Common Ground', *Linguistics and Philosophy*, 25 (5 - 6): 701 - 21.
- Williams, P. (1991). *The Alchemy of Race and Rights: A Diary of a Law Professor*, Cambridge: Harvard University Press.
- Wylie, A. (2004). 'Why Standpoint Matters', in *The Feminist Standpoint Theory Reader*, ed. S. Harding, ed. New York: Routledge.
- Zagzebski, L. (2013). 'A Defense of Epistemic Authority', *Res Philosophica*, 90 (2): 293 - 306.

Chapter 6 References

- Alcoff, L. (2005). *Visible Identities: Race, Gender, and the Self*, Oxford: Oxford University Press.
- Alfano, M. (2015). 'Becoming less unreasonable: A reply to Sherman', *Social Epistemology Review and Reply Collective*, 4 (7): 59 - 62.
- Anderson, D.E. (2017). 'Conceptual Competence Injustice', *Social Epistemology*, 31: 2, 210 - 223.

- Anderson, D.E. (2018). 'Yes, There Is Such a Thing as Conceptual Competence Injustice', *Social Epistemology Review and Reply Collective*, 7 (1): 26 - 35.
- Anderson, E. (2012). 'Epistemic Justice as a Virtue of Social Institutions', *Social Epistemology*, 31 (2): 163 - 173.
- Antony, L. (1995). 'Sisters Please, I'd Rather do it Myself', in *Philosophical Topics*, 23 (2): 59 - 94.
- Austin, J. (1946). 'Other Minds', *Proceedings of the Aristotelian Society*, Supplementary Volumes, 20: 148 - 187.
- Austin, J. (1962). *How to Do Things with Words*, London: Oxford University Press.
- Blome-Tillmann, M. (2009). 'Knowledge and Presuppositions', *Mind*, 118 (470): 241 - 94.
- Blome-Tillmann, M. (2014). *Knowledge and Presuppositions*, Oxford: Oxford University Press.
- Brister, E. (2009). 'Feminist Epistemology, Contextualism, and Philosophical Skepticism', *Metaphilosophy*, 40 (5): 671 - 688.
- Brister, E. (2017). 'Feminism and Contextualism', in *The Routledge Handbook to Epistemic Contextualism*, ed. J.J. Ichikawa. Taylor & Francis.
- Craig, E. (1990). *Knowledge in the State of Nature*, Oxford: Oxford University Press.
- Davies, A. (2016). 'How to Silence Content with Porn, Context and Loaded Questions', *European Journal of Philosophy*, 24 (2): 498 - 522.
- Dembroff, R. & Whitcomb, D. (forthcoming). 'Content-Focused Epistemic Injustice', *Oxford Studies in Epistemology*.
- DeRose, K. (1992). 'Contextualism and Knowledge Ascriptions', *Philosophy and Phenomenological Research*, 52: 931 - 29.
- DeRose, K. (2009). *The Case for Contextualism: Knowledge, Scepticism, and Context*, Oxford: Oxford University Press.

- Dotson, K. (2011). 'Tracking Epistemic Violence, Tracking Practices of Silencing', *Hypatia*, 26 (2): 236 - 257.
- Dotson, K. (2012). 'A Cautionary Tale: On Limiting Epistemic Oppression', *Frontiers: A Journal of Women Studies*, 33 (1): 24 - 47.
- Dotson, K. (2018). 'Accumulating Epistemic Power: A Problem with Epistemology', *Philosophical Topics*, 46 (1): 129 - 154.
- Dretske, F. (1970). 'Epistemic Operators', *Journal of Philosophy*, 67 (24): 1007 - 23.
- Ferguson, C. & Malouff, J. (2016). 'Assessing Police Classifications of Sexual Assault Reports: A Meta-Analysis of False Reporting Rates', *Archives of Sexual Behaviour*, 45: 1185 - 93.
- Fricker, M. (2007). *Epistemic Injustice: Power and the Ethics of Knowing*, Oxford: Oxford University Press.
- Gardiner, G. (2020). 'Relevance and Risk: How the Relevant Alternatives Framework Models the Epistemology of Risk', *Synthese*.
- Gardiner, G. (2019). 'The Reasonable and the Relevant: Legal Standards of Proof', *Philosophy & Public Affairs*, 47 (3): 288 - 318.
- Gerken, M. (2019). 'Pragmatic Encroachment and the Challenge from Epistemic Injustice', *Philosopher's Imprint*, 19 (15): 1 - 19.
- Goldman, A. (1976). 'Discrimination and Perceptual Knowledge', *Journal of Philosophy*, 73: 771 - 791.
- Haslanger, S. (2017). 'Culture and Critique', *Aristotelian Society Supplementary Volume XCI*: 149 - 73.
- Hawthorne, J. & Stanley, J. (2008). 'Knowledge and Action', *Journal of Philosophy*, 105 (10): 571 - 590.
- Hesni, S. (2018). 'Illocutionary Frustration', *Mind*, 127 (508): 947 - 976.

- Hookway, C. (2010). 'Some Varieties of Epistemic Injustice: Reflections on Fricker', *Episteme*, 7 (2): 151 - 163.
- Ichikawa, J.J. (2020). 'Contextual Injustice', *Kennedy Institute of Ethics Journal*, 30 (1): 1 - 30.
- Ichikawa, J.J. (2017). *Contextualising Knowledge: Epistemology and Semantics*, Oxford: Oxford University Press.
- Ichikawa, J.J. (2015). 'Ignorance and Presuppositions', *Mind*, 124 (496): 1207 - 1219.
- Jones, K. (2002). 'The Politics of Credibility', in *A Mind of One's Own: Feminist Essays on Reason and Objectivity*, ed. Louise Antony & Charlotte Witt. Westview Press.
- Jones, K. (2012). 'The Politics of Intellectual Self-Trust', *Social Epistemology*, 26 (2): 237 - 251.
- Kelp, C. (2011). 'What's the Point of 'Knowledge' Anyway?', *Episteme*, 8 (1): 53 - 66.
- Lai, C.K. et al. (2016). 'Reducing Implicit Racial Preferences: II. Intervention Effectiveness Across Time', *Journal of Experimental Psychology: General*, 145 (8): 1001 - 1016.
- Langton, R. (1993). 'Speech Acts and Unspeakable Acts', *Philosophy and Public Affairs*, 22 (4): 293 - 330.
- Langton, R. & Hornsby, J. (1998). 'Free Speech & Illocution', *Legal Theory*, 4 (1): 21 - 37.
- Langton, R. & West, C. (1999). 'Scorekeeping in a Pornographic Language Game', *Australasian Journal of Philosophy*, 77 (3): 303 - 319.
- Langton, R. (2018). 'Blocking as Counter-Speech', in *New Work on Speech Acts*, ed. D.Harris, D.Fogal, and M.Moss. New York: Oxford University Press.
- Lewis, D. (1996). 'Elusive Knowledge', *Australasian Journal of Philosophy*, 74: 549 - 567.
- Maitra, I. (2009). 'Silencing Speech', *Canadian Journal of Philosophy*, 39 (2): 309 - 338.
- Maitra, I. (2010). 'The Nature of Epistemic Injustice', *Philosophical Books*, 51 (4): 195 - 211.
- Manne, K. (2017). *Down Girl*, Oxford: Oxford University Press.

- McGowan, M.K. (2009). 'Oppressive Speech', *Australasian Journal of Philosophy*, 87 (3): 389 – 407.
- McKenna, R. (2017). 'Conversational Kinematics', in *The Routledge Handbook to Epistemic Contextualism*, ed. J.J. Ichikawa. Taylor & Francis.
- Medina, J. (2011). 'The Relevance of Credibility Excess in a Proportional View of Epistemic Injustice: Differential Epistemic Authority and the Social Imaginary', *Social Epistemology*, 25 (1): 15 – 35.
- Medina, J. (2013). *The Epistemology of Resistance*, Oxford: Oxford University Press.
- Pohlhaus, G. (2012). 'Relational Knowing and Epistemic Injustice: Towards a Theory of Willful Hermeneutical Ignorance', *Hypatia*, 27 (4): 715 – 735.
- Pohlhaus, G. (2014). 'Discerning the Primary Harm in cases of Testimonial Injustice', *Social Epistemology*, 28 (2): 99 – 114.
- Riggs, W. (2019). 'Open-Mindedness', in H. Battaly (Ed.), *The Routledge Handbook to Virtue Epistemology*, (pp.141 – 154). London: Routledge.
- Ross, L. (2020). 'The Virtue of Curiosity', *Episteme*, 17 (1): 105 – 120.
- Sherwood, H. (2019). 'Sexism, Vandalism, and Bullying: Inside the Boris Johnson-era Bullingdon Club', *The Observer*, 7th July 2019. Online at: <https://www.theguardian.com/politics/2019/jul/07/oxford-bullingdon-club-boris-johnson-sexism-violence-bullying-culture>. Accessed 16th June 2020.
- Stine, G. (1976). 'Scepticism, Relevant Alternatives, and Deductive Closure', *Philosophical Studies*, 29: 249 – 261.
- Turkheimer, D. (2017). 'Incredible Women: Sexual Violence and the Credibility Discount', *Pennsylvania Law Review*, 166: 1 – 58.
- Williams, M. (2001). 'Contextualism, Externalism, and Epistemic Standards', *Philosophical Studies*,

103 (1): 1 -23.

Williamson, T. (2000). *Knowledge and Its Limits*, Oxford: Oxford University Press.

Wylie, A. (2004). 'Why Standpoint Matters', in *The Feminist Standpoint Theory Reader*, ed. S. Harding, New York: Routledge.

Chapter 6.5. References

Blome-Tillmann, M. (2009). 'Knowledge and Presuppositions', *Mind*, 118 (470): 241 - 94.

Blome-Tillmann, M. (2013). 'Contextualism and the Knowledge Norms', *Pacific Philosophical Quarterly*, 94: 89 - 100.

Blome-Tillmann, M. (2014). *Knowledge and Presuppositions*, Oxford: Oxford University Press.

Cohen, S. (1986). 'Knowledge and Context', *Journal of Philosophy*, 83 (10): 574 - 583

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

Cohen, S. (1998). 'Contextualist Solutions to Epistemological Problems: Skepticism, Gettier, and the Lottery', *Australasian Journal of Philosophy*, 76 (2): 289 - 306.

Cohen, S. (1999). 'Contextualism, Skepticism, and the Structure of Reasons', *Philosophical Perspectives*, 13 (s13): 57 - 89.

DeRose, K. (1992). 'Contextualism and Knowledge Attributions', *Philosophy and Phenomenological Research*, 52: 913 - 929.

DeRose, K. (1995). 'Solving the Skeptical Problem', *Philosophical Review*, 104: 1 - 52.

DeRose, K. (2004). 'Single Scoreboard Semantics', *Philosophical Studies*, 119 (1 - 2): 1 -21.

DeRose, K. (2009) *The Case for Contextualism*, Oxford: Clarendon.

Heller, M. (1999). 'The Proper Role for Contextualism in Anti-Luck Epistemology', *Philosophical*

Perspectives, 33: 115 - 129.

Ichikawa, J.J. (2020). 'Contextual Injustice', *Kennedy Institute of Ethics Journal*, 30 (1): 1 - 30.

Ichikawa, J.J. (2017). *Contextualising Knowledge: Epistemology and Semantics*, Oxford: Oxford University Press.

Lewis, D. (1996). 'Elusive Knowledge', *Australasian Journal of Philosophy*, 74 (4): 549 - 67.

McKenna, R. (2017). 'Conversational Kinematics', in *Routledge Handbook of Epistemic Contextualism*, ed. J.J. Ichikawa. Taylor & Francis (pp. 321 - 331).

Neta, R. (2002). 'S Knows that P', *Nous*, 36 (4): 663 - 681.

Neta, R. (2004). 'Contextualism and the Problem of the External World', *Philosophy and Phenomenological Research*, 66 (1): 1 - 31.

Rieber, S. (1998). 'Skepticism and Contrastive Explanation', *Nous*, 35 (4): 477 - 451.

Rysiew, P. (2021) 'Epistemic Contextualism', *The Stanford Encyclopedia of Philosophy* (Spring 2021 Edition), Edward N. Zalta (ed.), URL
=<<https://plato.stanford.edu/archives/spr2021/entries/contextualism-epistemology/>>.

Stine, G. (1976). 'Skepticism, Relevant Alternatives, and Deductive Closure', *Philosophical Studies*, 29 (4): 249 - 261.

Chapter 7 References

Baehr, J. (2013). 'Educating for Intellectual Virtues: From Theory to Practice', *Journal of Philosophy of Education: From Theory to Practice*, 47 (2): 248 - 262.

Begby, E. (2021). *Prejudice: A Study in Non-ideal Epistemology*, Oxford: Oxford University Press.

Blome-Tillmann, M. (2014). *Knowledge and Presuppositions*, Oxford: Oxford University Press.

- Cassam, Q. (2015). 'Stealthy Vices', *Social Epistemology Review and Reply Collective*, 4 (10): 19 – 25.
- Cassam, Q. (2016). 'Vice Epistemology', *The Monist*, 99 (2): 159 – 180.
- Cohen, S. (1988). 'How to be a Fallibilist', *Philosophical Perspectives*, 2: 91 – 123.
- Crerar, C. (2017). 'Motivational Approaches to Intellectual Virtue', *Australasian Journal of Philosophy*, 96 (4), 753 – 766.
- DeRose, K. (1995). 'Solving the Skeptical Problem', *Philosophical Review*, 104 (1): 1 – 52.
- DeRose, K. (2004). 'Single Scoreboard Semantics', *Philosophical Studies*, 119 (1 – 2): 1 – 21.
- Diaz-Leon, E. (2017). 'Epistemic Contextualism and Conceptual Ethics'. In J.J. Ichikawa (Ed.), *The Routledge Handbook to Epistemic Contextualism*, (pp. 71 – 81). London: Routledge.
- Fricker, M. (2007). *Epistemic Injustice: Power and the Ethics of Knowing*, Oxford: Oxford University Press.
- Greco, J. (2002). 'Virtues in Epistemology', in *The Oxford Handbook of Epistemology*, Ed. Paul, K. Moser, (pp.287 – 315). New York: Oxford University Press.
- Greco, J. (2004). 'A Different Sort of Contextualism', *Erkenntnis*, 61: 383 – 400.
- Greco, J. (2008). 'What's Wrong with Contextualism?', *The Philosophical Quarterly*, 58 (232): 416 – 436.
- Grice, P. (1989). *Studies in the Way of Words*, Cambridge MA: Harvard University Press.
- Hursthouse, R. (1991). 'Virtue Theory and Abortion', *Philosophy & Public Affairs*, 20: 223 – 246
- Kidd, I.J. (2016). 'Charging Others with Epistemic Vice', *The Monist*, 99 (3): 181 – 197.
- Lewis, D. (1996). 'Elusive Knowledge', *Australasian Journal of Philosophy*, 74 (4): 549 – 567.
- Mill, J.S. (1977 [1859]). *On Liberty, Collected Works XVII*, Toronto: University of Toronto Press.

- Montmarquet, J. (2019). 'Do Epistemic Virtues Require a Motivation for Truth?', In H. Battaly (Ed.), *The Routledge Handbook to Virtue Epistemology*, (pp. 37 - 46). London: Routledge.
- Montmarquet, J. (1993). *Epistemic Virtue and Doxastic Responsibility*, Lanham, MD: Rowman & Littlefield.
- Plunkett, D. & Sundell, T. (2013). 'Disagreement and the Semantics of Normative and Evaluative Terms', *Philosophers Imprint*, 13 (23): 1 - 37.
- Riggs, W. (2019). 'Open-Mindedness', In H. Battaly (Ed.), *The Routledge Handbook to Virtue Epistemology*, (pp. 141 - 154). London: Routledge.
- Roberts, R. & Wood, J. (2007). *Intellectual Virtues*, Oxford: Oxford University Press.
- Sherman, B. (2016). 'There's No Justice: Why Pursuit of a Virtue is Not the Solution to Epistemic Injustice', *Social Epistemology*, 30 (3): 229 - 250.
- Slote, M. (1995). 'Agent-Based Virtue Ethics', *Midwest Studies in Philosophy*, 20: 83 - 101.
- Sosa, E. (1991). *Knowledge in Perspective*, New York, NY: Cambridge.
- Swanton, C. (2001). 'A Virtue Ethical Account of Right Action', *Ethics*, 112: 32 - 52.
- Zagzebski, L. (1996). *Virtues of the Mind: An Inquiry into the Nature of Virtue and the Ethical Foundations of Knowledge*, New York, NY: Cambridge University Press.

Chapter 8 References

- Basu, R. (2019). 'Radical Moral Encroachment: The Moral Stakes of Racist Beliefs', *Philosophical Issues*, 29(1): 9 - 23.
- Basu, R. & Schroeder, M. (2019). 'Doxastic Wronging', in Kim, B. & McGrath, M. (eds.), *Pragmatic Encroachment in Epistemology*, pp.181 - 205. Routledge.
- Blome-Tillmann, M. (2009). 'Knowledge and Presuppositions', *Mind*, 118 (470): 241 - 94.

- Blome-Tillmann, M. (2014). *Knowledge and Presuppositions*, Oxford: Oxford University Press.
- Bolinger, R. (2020). 'The Rational Impermissibility of Accepting (Some) Racial Generalizations', *Synthese*, 197: 2415 - 2431.
- Bolinger, R. (forthcoming). 'Varieties of Moral Encroachment', *Philosophical Perspectives*.
- Cohen, S. (1999). 'Contextualism, Skepticism, and the Structure of Reasons', *Philosophical Perspectives*, 13 (s13): 57 - 89.
- Crewe, B. & Ichikawa, J.J. (forthcoming). 'Rape Culture and Epistemology', in *Applied Epistemology*, ed. Lackey, J. Oxford: Oxford University Press.
- DeRose, K. (2004). 'Single Scoreboard Semantics', *Philosophical Studies*, 119 (1 - 2): 1 -21.
- DeRose, K. (2005). 'The Ordinary Language Basis for Contextualism, and the New Invariantism', *The Philosophical Quarterly*, 55 (219) 172 - 198.
- DeRose, K. (2009) *The Case for Contextualism*, Oxford: Clarendon.
- Fricker, M. (2007). *Epistemic Injustice: Power and the Ethics of Knowing*, Oxford: Oxford University Press.
- Fritz, J. (2019). 'Moral Encroachment and Reasons of the Wrong Kind', *Philosophical Studies*, Online first.
- Fritz, J. (2017). 'Pragmatic and Moral Encroachment', *Pacific Philosophical Quarterly*, 98 (s1): 643 - 661.
- Gendler, (2011). 'On the Epistemic Costs of Implicit Bias', *Philosophical Studies*, 156 (1): 33 - 63.
- Gertken, J. & Kiesewetter, B. (forthcoming). 'The Right and Wrong Kind of Reasons', *Philosophy Compass*.
- Guerrero, A. (2007). 'Don't Know, Don't Kill: Moral Ignorance, Culpability, and Caution', *Philosophical Studies*, 136: 59 - 97.
- Hannon, M. (2020). 'Why Purists Should be Infallibilists', *Philosophical Studies*, 177: 689 - 704.

- Ichikawa, J.J. (2017). *Contextualising Knowledge: Epistemology and Semantics*, Oxford: Oxford University Press.
- Lewis, D. (1996). 'Elusive Knowledge', *Australasian Journal of Philosophy*, 74 (4): 549 – 67.
- McKenna, R. (2011). 'Interests Contextualism', *Philosophia*, 39: 741 – 750.
- Moss, S. (2018). 'Moral Encroachment', *Proceedings of the Aristotelian Society*, 118 (2): 177 – 205.
- Pace, M. (2011). 'The Epistemic Value of Moral Considerations: Justification, Moral Encroachment, and James' 'Will to Believe'', *Nous*, 45 (2): 239 – 268
- Schroeder, M. (2018). 'When Beliefs Wrong', *Philosophical Topics*, 46 (1): 115 – 127.
- Schaffer, J. (2005). 'What Shifts? Thresholds, Standards, or Alternatives?', in Preyer, G. & Peter, G. (eds.), *Contextualism in Philosophy: Knowledge, Meaning, and Truth*. Oxford: Oxford University Press.
- Stroud, S. (2006). 'Epistemic Partiality in Friendship', *Ethics*, 116 (3): 498 – 524.

