


**ON THE CONTRARY
DISAGREEMENT, CONTEXT, AND RELATIVE TRUTH**

Torfinn Thomesen Huvenes

**A Thesis Submitted for the Degree of PhD
at the
University of St. Andrews**



2012

**Full metadata for this item is available in
Research@StAndrews:FullText
at:**

<http://research-repository.st-andrews.ac.uk/>

Please use this identifier to cite or link to this item:

<http://hdl.handle.net/10023/3360>

This item is protected by original copyright

On the Contrary

Disagreement, Context, and Relative Truth

Torfinn Thomesen Huvenes

A PhD thesis to be submitted to the University of Oslo
for a double badge degree from
the University of Oslo and the University of St Andrews

Department of Philosophy, Classics, History of Art and Ideas
University of Oslo

Department of Philosophy
School of Philosophical, Anthropological and Film Studies
University of St Andrews

February 2011

1. Candidate's declarations:

I, Torfinn Thomesen Huvenes, hereby certify that this thesis, which is approximately 57,500 words in length, has been written by me, that it is the record of work carried out by me and that it has not been submitted in any previous application for a higher degree.

I was admitted as a research student in 2007 and as a candidate for the degree of Doctor of Philosophy in 2007; the higher study for which this is a record was carried out in the University of St Andrews between 2007 and 2011.

Date 20/01/12 signature of candidate

2. Supervisor's declaration:

I hereby certify that the candidate has fulfilled the conditions of the Resolution and Regulations appropriate for the degree of Doctor of Philosophy in the University of St Andrews and that the candidate is qualified to submit this thesis in application for that degree.

Date 05/04/12 signature of supervisor

3. Permission for electronic publication: *(to be signed by both candidate and supervisor)*

In submitting this thesis to the University of St Andrews we understand that we are giving permission for it to be made available for use in accordance with the regulations of the University Library for the time being in force, subject to any copyright vested in the work not being affected thereby. We also understand that the title and the abstract will be published, and that a copy of the work may be made and supplied to any bona fide library or research worker, that my thesis will be electronically accessible for personal or research use unless exempt by award of an embargo as requested below, and that the library has the right to migrate my thesis into new electronic forms as required to ensure continued access to the thesis. We have obtained any third-party copyright permissions that may be required in order to allow such access and migration, or have requested the appropriate embargo below.

The following is an agreed request by candidate and supervisor regarding the electronic publication of this thesis:

Access to printed copy and electronic publication of thesis through the University of St Andrews.

Date 05/04/12 signature of candidate signature of supervisor

A supporting statement for a request for an embargo must be included with the submissions of the draft copy of the thesis. Where part of a thesis is to be embargoed, please specify the part and the reasons.

Acknowledgements

My greatest thanks have to go to my main supervisor, Professor Herman Cappelen. Throughout countless meetings, skype conversations, and emails, he has provided invaluable advice, comments, and support. There is no way to overstate how important his continuing support has been. Indeed, it is hard to imagine a better supervisor.

I also owe a large debt to my second supervisor, Professor Jonathan Schaffer. His comments and support have been tremendously important. I have been very lucky to have two great supervisors.

While I have been working on my thesis, I have been splitting my time between St Andrews and Oslo. In St Andrews, I have been a member of the Arché Research Centre at the University of St Andrews. In Oslo, I have been a member of the Centre for the Study of Mind in Nature (CSMN) at the University of Oslo. I am extremely grateful for having been given the opportunity to be a part of these stimulating research environments. The members of these research centres deserve great thanks for contributing to these great research environments. I have also benefited a lot from interactions with many of the visitors at Arché and CSMN, as well as participants at conferences, workshops, seminars, and reading groups that I have attended. The list is endless, but it includes Nicholas Allott, Ralf Bader, Björn Brodowski, Jessica Brown, Einar Bøhn, Colin Caret, Andreas Brekke Carlsson, Robyn Carston, Yuri Cath, Timothy Chan, Josh Clarkson, Stewart Cohen, Michael De, Laura Delgado, Josh Dever, Dylan Dodd, Janice Dowell, Douglas Edwards, Andy Egan, Christel Fricke, Olav Gjelsvik, Patrick Greenough, Eline Busck Gundersen, Carsten Hansen, Katherine Hawley, John Hawthorne, Richard Heck, Sanna Hirvonen, Ole Hjortland, Thomas Hodgson, Heine Holmen, Jonathan Ichikawa, Ben Jarvis, Carrie Jenkins, Dirk Kindermann,

Ira Kiourti, Georg Kjöll, Max Kölbel, Ernie Lepore, Julia Langkau, Peter Ludlow, Federico Luzzi, John MacFarlane, Ishani Maitra, Julien Murzi, Anders Nes, Andrea Onofri, Walter Pedrialli, Laura Porro, Graham Priest, Simon Prosser, Brian Rabern, Bjørn Ramberg, Stephen Read, Francois Recanati, Kari Refsdal, Mathias Sagdahl, Kalle Sandås, Susanna Schellenberg, Kristin Selmer, Stewart Shapiro, Knut Olav Skarsaune, Martin Smith, Jason Stanley, Rachel Sterken, Hallvard Stette, Isidora Stojanovic, Anders Strand, Margot Strohminger, Paula Sweeney, Chiara Tabet, Brian Weatherson, Deirdre Wilson, Crispin Wright, Elia Zardini, Dan Zeman, and Tore Øgaard. I apologise for anyone who might have been left out.

I would like to extend a special thanks to Derek Ball, Paul Dimmock, Guðmundur Andri Hjálmarsson, Dilip Ninan, Anders Schoubye, Daniele Sgaravatti, and Andreas Stokke. They have influenced my thinking about a wide range of philosophical topics, and I cannot exaggerate how much I have enjoyed the countless discussions and conversations.

Finally, I want to thank my family and those of my friends who I have not already mentioned. Without their support I would have retained even less of my sanity, and this thesis would not have been a reality. In particular, I would like to thank my parents, Rune Huvenes and Siri Thomesen.

Contents

Introduction	9
1 Disagreement and Relative Truth	15
1.1 Relative Truth	16
1.2 The Argument from Disagreement	18
1.3 A Problem for the Simple Relativist Picture	21
1.3.1 Worlds Instead of Times	23
1.3.2 Other Parameters	27
1.4 Two Relativist Replies	30
1.4.1 The First Reply	30
1.4.2 The Second Reply	34
1.5 Indexical Relativism and Nonindexical Contextualism	40
2 Agreement and Disagreement	45
2.1 The Questions	47
2.2 The Doxastic View	51
2.3 A Problem for the Doxastic View	55
2.4 Further Problems for the Doxastic View	60
2.5 Relativism to the Rescue?	63
2.6 Disagreement in Attitude	65
2.7 Conflicting Attitudes	69
2.8 Faultless Disagreement	74
3 Predicates of Taste	79
3.1 Contextualism and Relativism	80
3.2 The Problem of Disagreement	82

3.3	More Disagreement Data	83
3.4	Doubts about the Data	87
3.5	Other Disagreement Markers	90
3.6	Conflicting Attitudes	95
4	Epistemic Modals	99
4.1	Contextualism and Relativism	101
4.2	The Data	103
4.3	More Data	106
4.4	Psychological Bias	108
4.5	More on Disagreement and Retraction	110
4.6	Stephenson on Disagreement	115
4.7	The Disagreement Constraint	117
4.8	The Retraction Constraint	120
4.9	Disagreement and Credences	123
4.10	The Pragmatics of Epistemic Modals	127
	4.10.1 Proffering	130
	4.10.2 The Common Propositional Space	131
	4.10.3 Doxastic Advice	132
4.11	Further Problems	134
5	Knowledge Ascriptions and Indirect Speech Reports	139
5.1	Contextualism about Knowledge Ascriptions	141
5.2	Indirect Speech Reports	143
5.3	Nonindexical Contextualism	144
5.4	Mixed Quotation	147
5.5	Salience and Parasites	149
5.6	Semantics and Propositional Content	152
5.7	The Operator Account	153
5.8	Nonindexical Contextualism After All?	156
5.9	Lambda Abstraction	158
5.10	Taking Stock and Looking Ahead	163
5.11	Monsters	166
5.12	Another Problem	168

Conclusion	175
References	179

Introduction

The Problem

What does it take for two individuals to disagree? That is a question that has become relevant for the debate about contextualism and relativism. Proponents of relativism have argued that, for a wide range of expressions, relativism does better than contextualism when it comes to giving an account of disagreement.¹ There are cases in which contextualism incorrectly predicts that the speakers do not disagree. Relativism is supposed to do better by delivering the prediction that the speakers disagree.

The idea behind contextualism is familiar. The claim is that a sentence is context-dependent in the sense that it express different propositions in different contexts. This is the standard way to think about sentences that contain indexicals like ‘I’ and ‘here’. A sentence like (1) expresses different propositions in different contexts depending on who the agent or speaker of the context is.²

(1) I am a student.

Contextualism, as a view about a certain class of expressions, is the view that sentences that contain the relevant expressions express different propositions

¹Relativism has been proposed as a view of, among other things, deontic modals, epistemic modals, future contingents, gradable adjectives, knowledge ascriptions, and predicates of taste. See e.g. Kölbel (2002, 2004), MacFarlane (2003, 2005a, 2007, 2008, forthcoming-a, forthcoming-b), Richard (2004, 2008), Egan, Hawthorne, and Weatherson (2005), Lasersohn (2005), Egan (2007, 2010), Stephenson (2007), and Kolodny and MacFarlane (2010).

²See Kaplan (1989) for a classic treatment of indexicals along these lines.

in different contexts.³ For instance, let us consider a contextualist treatment of ‘tasty’. The relevant view is that sentences like (2) express different propositions in different contexts, depending on whose taste that is relevant.

(2) Haggis is tasty.

This treatment of ‘tasty’ has come under criticism.⁴ The problem is that someone who sincerely utters (2) and someone who sincerely utters (3) seem to disagree.

(3) Haggis is not tasty.

A contextualist treatment of ‘tasty’ is supposed to be unable to account for this. If different standards of taste are in play, the relevant sentences express different propositions. But then the worry is that the situation is just like a situation in which someone sincerely utters (1), and someone else sincerely utters (4).

(4) I am not a student.

There is no pressure to say that the speakers disagree in this case. The speakers just believe different and compatible propositions.

According to contextualism, sentences have different truth-values relative to different contexts in virtue of expressing different propositions in different contexts. Proponents of relativism have suggested that we can do better if we take truth to be relative in a more interesting sense. For instance, we can think of sentences like (2), not as expressing different propositions in different contexts, but as expressing propositions that vary in truth-value across individuals or standards of taste. This view is supposed to do better than contextualism when it comes to respecting the judgement that the speaker of (2) and the speaker of (3) disagree. In particular, it allows us to say that there is a proposition such that one party believes that proposition, and the other party believes its negation.

³This is not meant to include any context-dependence not associated with the relevant expression, nor does it include any context-dependence associated with features like tense. There will typically be a fairly specific sense in which the relevant expressions are supposed to be context-dependent.

⁴See e.g. Kölbel (2002, 2004), Lasersohn (2005), MacFarlane (2007), and Stephenson (2007).

The Plan

This challenge serves as the starting point for my discussion of the issues concerning contextualism, relativism, and disagreement. Focusing on predicates of taste and epistemic modals, I argue that the argument against contextualism is problematic.⁵ In order to appreciate why the argument is problematic, we need to think about what it takes for two individuals to agree or disagree.⁶ A recurring theme in the discussion is that we should not assume a too restrictive view of agreement and disagreement.

I argue that agreement and disagreement should not always be understood in terms of what the parties believe. Whether two individuals agree or disagree can depend on what they prefer, what they like, what they want, and so forth. Apart from being philosophically interesting in their own right, these issues are relevant for the debate about contextualism and relativism. I argue that if we reject the view that disagreement is always a matter of there being a proposition such that one party believes that proposition and the other party believes its negation, there is more room for contextualists to explain the allegedly problematic cases of disagreement.

My claim is that, at least in the case of predicates of taste and epistemic modals, the disagreement data do not give us a reason for preferring relativism over contextualism. For the most part, I focus on defending contextualism rather than criticising relativism.⁷ The goal is to undermine the motivation for rela-

⁵There are several reasons for focusing on predicates of taste and epistemic modals. These expressions have played a central role in the debate between contextualists and relativists. The views in question are therefore fairly well-developed. A relativist treatment of predicates of taste has been defended by Kölbel (2002, 2004), Lasersohn (2005), MacFarlane (2007), and Stephenson (2007). A relativist treatment of epistemic modals has been defended by e.g. Egan et al. (2005), Egan (2007), Stephenson (2007), and MacFarlane (forthcoming-a). Predicates of taste and epistemic modals are perhaps also less likely to be associated with long-standing philosophical debates in aesthetics, epistemology, and moral philosophy, than, say, knowledge ascriptions or moral expressions.

⁶The argument might be problematic in more than one way. See e.g. Glanzberg (2007), Cappelen and Hawthorne (2009), and Schaffer (forthcoming-a) for relevant discussion.

⁷In an important sense the focus is on contextualism rather than relativism. While relativism allows us to bring the problem of disagreement into focus, the problem of disagreement is not essentially tied to relativism. There are other alternatives to contextualism and relativism that could have played a more important role in the discussion. For instance, I could have paid more attention

tivism by showing that contextualism can also account for the relevant cases of disagreement. In light of the special role that considerations involving disagreement have played in the debate, I take it that it would be highly significant if it turned out that these considerations do not motivate relativism after all.⁸

It is still important to recognise that the problem of disagreement is not the only problem for contextualism. It has also been argued that contextualist views have problems with attitude and indirect speech reports.⁹ While I do not intend to address every problem on behalf of contextualism, I think it is worth saying something about how contextualists can deal with indirect speech reports. I argue that this problem, unlike the problem of disagreement, requires the contextualist semantics to be modified in certain respects. However, this can be done without violating the central commitments of contextualism as long as we pay attention to the semantic implementation of contextualism.

Overview of the Chapters

What follows is a brief overview of the chapters. While the chapters are connected in terms of the issues that they deal with, it is meant to be possible to read them independently. They are supposed to stand on their own as individual essays. That means that there is a certain amount of overlap and repetition.

Chapter 1 mainly serves as a more detailed introduction to the issues concerning relativism and disagreement. I examine the basis for thinking that relativism is in a better position to account for disagreement than contextualism. This involves getting clear on how relativists ought to think about disagreement. I discuss an argument against a simple relativist account of disagreement. The argument is

to invariantist and expressivist views.

⁸I am interested in relativist views insofar as they are motivated by considerations involving agreement or disagreement. There are ways of arguing that propositional truth is relative that have nothing to do with agreement or disagreement. This has recently been emphasised by Weatherson (2011). I am not in the business of arguing against someone who is impressed by Kaplan's (1989) so-called 'operator argument' or Lewis' (1979) account of self-locating or *de se* attitudes and thinks that is a reason for taking propositional truth to be relative. See e.g. Cappelen and Hawthorne (2009) for a more general criticism of attempts to argue that propositional truth is relative.

⁹See e.g. Hawthorne (2004), Cappelen and Lepore (2005), Brogaard (2008), and MacFarlane (2009).

based on observations about what happens to disagreement if the relevant propositions are true or false relative to parameters like worlds and times. I consider two lines of response on behalf of relativism. The first response involves denying that propositions are true or false relative to the parameters in question. The second response amounts to the claim that the parameters that relativists want to posit are different from parameters like worlds and times. Both lines of response are viable in the sense that they allow relativists to avoid the conclusion that relativism is no better off than contextualism when it comes to dealing with disagreement. However, they are associated with different commitments.

Chapter 2 deals with more general questions about agreement and disagreement. I argue against a view I call the ‘doxastic view’ of agreement and disagreement. That is the view that whether two parties agree or disagree is always a matter of what they believe. I argue that there are cases of disagreement that cannot be adequately handled by the doxastic view. Instead I develop a more pluralist view of disagreement, taking Stevenson’s (1937, 1944, 1963) notion of ‘disagreement in attitude’ as my starting point. According to this view, agreement and disagreement can involve a wide range of attitudes. That is more or less how expressivists like Blackburn (1984, 1998) and Gibbard (1990, 2003) have thought about agreement and disagreement, but I argue this way of thinking about agreement and disagreement is not essentially connected with expressivism.

Chapter 3 is a discussion of contextualism and relativism about predicates of taste.¹⁰ There are allegedly cases in which contextualism about predicates of taste incorrectly predicts that the speakers do not disagree. This argument turns out to rest on a problematic conception of disagreement. I argue that there are other cases of disagreement, cases in which speakers are talking about what they like, that we ought to take into account. If these cases are taken into account, the original disagreement data are less surprising from a contextualist point of view. Furthermore, I suggest that we can make sense of all of these cases if they are understood as involving a conflict of non-doxastic attitudes. In other words, the idea is to adopt something like the view of agreement and disagreement de-

¹⁰Chapter 3 is a revised version of a paper forthcoming in the *Australasian Journal of Philosophy*. See Huvenes (forthcoming).

veloped in the previous chapter.

Chapter 4 is about contextualist and relativist treatments of epistemic modals. Contextualism about epistemic modals allegedly has problems when it comes to dealing with disagreement and retraction data. Following a suggestion from von Fintel and Gillies (2008), I argue that the disagreement and retraction is targeting the embedded proposition, not the modal proposition expressed by the relevant sentence. Certain observations made by Stephenson (2007) and MacFarlane (forthcoming-a) might be seen as undermining this line of response. However, that turns out not to be the case. In particular, I argue that we ought to reject the constraint that disagreement and retraction can only target a proposition that is believed or asserted. I suggest that the basis for the disagreement is a difference in credences or degrees of belief, not a difference in outright beliefs. I also examine different ways of thinking about the pragmatics of epistemic modals as a means of providing a better understanding of the disagreement and retraction data. The idea is that sentences containing epistemic modals can be used to perform two distinct speech acts. While this idea is promising, the conclusion is that more research is needed.

Chapter 5 is different from the other chapters insofar as it is about indirect speech reports and not about disagreement. I discuss whether contextualism about knowledge ascriptions can deal with cases in which knowledge ascriptions are embedded in indirect speech reports. I argue that accounts that appeal to mixed quotation and salience are inadequate. Instead we need to make more substantive assumptions about the semantics of knowledge ascriptions. I present two accounts that do a better job of getting the right results. However, both accounts come with significant commitments about the semantic implementation of contextualism. The first account is based on the idea of adding an epistemic standard parameter to the index and treating ‘says’ as an operator that shifts the value of that parameter. The second account is based on the idea of positing a covert variable associated with ‘knows’ that can be bound by a lambda operator. While considerations involving indirect speech reports do not directly undermine contextualism, these considerations are still relevant for the overall assessment of contextualism.

Chapter 1

Disagreement and Relative Truth

Introduction

There has been a lot of recent interest in the idea that some form of relativism about truth can provide the basis for a successful semantics for a wide range of expressions. This includes, among other things, epistemic modals, knowledge ascriptions, and predicates of taste.¹ An important motivation for these relativist views is that they are supposed to be better suited for handling disagreement than some of their rivals. This raises a general question about what happens to disagreement if truth is relative. In particular, what happens to disagreement if truth is relative in the sense that relativists like Kölbel (2002), Lasersohn (2005), and MacFarlane (2005b, 2007) are talking about? The goal in this chapter is to discuss relativist views of disagreement.

One might be more or less ambitious on behalf of relativism. A fairly modest goal would be to show that relativists can respond to the objections that have been raised against relativist accounts of disagreement. A more ambitious goal would be to argue that the relativist picture is independently motivated. The following discussion will focus on the more modest goal. This does not mean that it is not important to be clear about the extent to which parts of the view are independently motivated or not. Among other things, I am interested in

¹See e.g. Kölbel (2002, 2004), Richard (2004, 2008), Egan et al. (2005), Lasersohn (2005), MacFarlane (2005b), Egan (2007, 2010) and Stephenson (2007).

whether relativists can use the connection between disagreement and truth as a means of motivating their view of disagreement. In any case, my goal is not to settle whether relativists have a plausible account of disagreement. The main purpose of this chapter is provide an introduction to some of the issues involving relativism and disagreement, not to argue for or against relativism.

In §1.1 I distinguish between different ways in which truth can be relative, and offer a preliminary characterisation of relativism. In §1.2 I talk about how considerations involving disagreement are supposed to motivate relativism. In §1.3 I discuss a problem for a simple relativist account of disagreement that is based on what happens to disagreement if propositions are true or false relative to parameters like worlds and times. In §1.4 I consider two lines of response on behalf of relativism. In §1.5 I briefly discuss two positions, indexical relativism and non-indexical contextualism, as potential alternatives to contextualism and relativism.

1.1 Relative Truth

In order to get clear on what a relativist view amounts to, there are two questions that one ought to consider. The first question is what the relevant bearers of truth and falsity are. In this context, it matters whether we are talking about the truth of sentences, utterances, or propositions. The view that sentences are true or false relative to contexts is not particularly controversial, and it does not amount to an interesting relativist thesis. There is nothing surprising about the idea that a sentence like (5) is true relative to some contexts, but false relative to others.

(5) I am hungry.

For the purpose of the following discussion I will assume that the relevant truth-bearers are propositions. This excludes certain views from the discussion. In particular, I am not going to focus on views according to which sentences express different propositions relative to different contexts of assessment or interpretation.² I will only discuss these views briefly in §1.5. Insofar as we are interested in questions about disagreement, this simplifies the discussion. If propositions are the objects of belief, and disagreement is a matter of what we believe, it makes

²See e.g. MacFarlane (2005b), Cappelen (2008a, 2008b), and Weatherson (2009) for relevant discussion.

sense to focus on propositional truth rather than, say, sentential truth or utterance truth.³

The second question is what truth is supposed to be relative to. Even if we are only interested in views on which propositional truth is relative, this does not guarantee that we end up with an interesting relativist view. For instance, while it is not universally accepted, many take propositions to be true or false relative to possible worlds.⁴ The view that there are tensed or temporally neutral propositions, propositions that are true or false relative to times, is somewhat more controversial.⁵ However, this still does not amount to anything like an interesting relativist view in the sense that I am interested in.

Stanley (2005) suggests that on a relativist view, propositional truth is relative to a ‘non-standard’ parameter. Without a clear idea of what counts as a ‘non-standard’ parameter, this is perhaps not very helpful, but it is not easy to improve on this characterisation. Relativists differ when it comes to what they take truth to be relative to. For instance, Kölbel (2002) takes relativism to be the view that propositional truth is relative to perspectives. However, MacFarlane (2005b) argues that what matters is whether propositional truth is relative to contexts of assessment.⁶ These differences will be relevant as the discussion proceeds, but for now I am going to be content with something like Stanley’s admittedly vague characterisation. In order to get a better grip on the relativist views that I am interested in, we need to look at the way they are motivated.

However, there is an issue that I would like to flag. It is natural to talk about contextualism and relativism as views about the semantics of a certain class of expressions, and I will continue to do so for the purpose of the following discussion. However, I will for the most part be talking about propositions as the

³See Chapter 2 for further relevant discussion.

⁴See e.g. Schaffer (forthcoming-b) for a defence of the view that propositions are not true or false relative to possible worlds. On his view, propositions are specific with respect to features like world and time. See e.g. also Cappelen and Hawthorne (2009).

⁵For criticism of the view that propositions are true or false relative to times, see e.g. Richard (1981), King (2003), and Cappelen and Hawthorne (2009). See e.g. Recanati (2007) for a recent defence of the view that there are contents that are true or false relative to times.

⁶Strictly speaking, MacFarlane (2005b) takes relativism to be the view that there is a sentence that has different truth-values relative to different contexts of assessment. However, he seems to prefer the version of the view he calls ‘propositional relativism’. According to this view, there are propositions that have different truth-values relative to different contexts of assessment.

objects of belief. As long as we assume that the propositions that serve as the objects of the relevant beliefs are the propositions expressed by the relevant sentences, these complications do not have to bother us too much. In any case, I will assume that we can move back and forth between the level of language and the level of thought without too much worry.⁷ While this assumption might prove to be naive, it simplifies the discussion.

1.2 The Argument from Disagreement

What the relativist views that I am interested in have in common, is that they are, at least to some extent, motivated by considerations involving disagreement. In particular, they are supposed to do a better job than more traditional contextualist views when it comes to dealing with disagreement. For the purpose of the following discussion, I will simply ignore views that are not motivated by disagreement in this way. In order to illustrate this, let us consider a relativist treatment of sentences containing expressions like ‘beautiful’.⁸ On this view, propositions are true or false relative to aesthetic standards. The basic idea is that a sentence like (6) expresses a proposition that can be true relative to one set of aesthetic standard, but false relative to another.

(6) Picasso’s *Guernica* is beautiful.

According to a more traditional contextualist treatment, a sentence like (6) would express different propositions in different contexts. For instance, if the sentence

⁷The assumption that we can move back and forth between the level of language and the level of thought in this way, also holds for later chapters. I am also assuming that we do not have to distinguish between what Lewis (1980) called the ‘semantic value’ of a sentence and what he called the ‘propositional content’. While the former is what is assigned to a sentence by a compositional semantic theory, the latter is what serves as the objects of speech acts and propositional attitudes. Lewis argued that these can come apart, but for the most part I will assume that we can identify the semantic value of a sentence with its propositional content. However, I will return to this assumption in Chapter 5.

⁸This is only meant to serve as an illustration, and I will not engage in a substantive discussion about the correct treatment of expressions like ‘beautiful’ either in this chapter or in later chapters. For the purpose of the following discussion I will also assume that the only views on the table are contextualism and relativism. This is obviously an enormous over-simplification, but it simplifies the discussion significantly and nothing important turns on it.

is uttered by John, it might express the proposition that Picasso's painting *Guernica* is beautiful by John's standards. However, if the sentence is uttered by Mary, it might express the proposition that the painting is beautiful by Mary's standards. There is a sense in which truth is relative on the contextualist picture, but it is the familiar sense in which sentences are true or false relative to contexts in virtue of expressing different propositions in different contexts. On the relativist picture, (6) expresses the proposition that Picasso's *Guernica* is beautiful, but that proposition may be true relative John's standards and false relative to Mary's standards.

A worry with the contextualist picture is that it fails to make the right predictions about when speakers disagree. Suppose that John sincerely asserts (6) and that Mary sincerely asserts (7).

(7) Picasso's *Guernica* is not beautiful.

It would seem that Mary and John disagree. However, a contextualist treatment is supposedly unable to respect this judgement. The problem is that it is not clear that there is a proposition such that one of them believes that proposition and the other believes its negation. If John believes the proposition that the painting is beautiful by his standards, and Mary believes the proposition that the painting is not beautiful by her standards, this does not seem to amount to a disagreement. While there are no doubt other ways of spelling out the argument against contextualism in more detail, this way of spelling out the argument seems to capture the basic idea:⁹

1. The parties disagree. (Ass.)
2. If the parties disagree, there is a proposition such that one party believes that proposition and the other believes its negation. (Ass.)
3. If contextualism is true, there is not a proposition such that one party believes that proposition and the other believes its negation. (Ass.)
4. There is a proposition such that one party believes that proposition and the other party believes its negation. (from 1 and 2 by modus ponens)

⁹I am presenting the argument against contextualism as a deductive argument and one might take issue with this, but as far as I can see, nothing substantial turns on that assumption here. See e.g. Weatherson (2011) for relevant discussion.

∴ 5. Contextualism is not true. (from 3 and 4 by modus tollens)

The first three premises are certainly not beyond question, and any of them may be challenged.¹⁰ However, at this point, I am not interested in discussing whether this amounts to a good argument against contextualism. The question is whether relativism can do better than contextualism on the assumption that the argument is sound. Insofar as the alleged problem with contextualism is that it does not predict that there is a proposition such that John believes that proposition and Mary believes its negation, one might think that relativism has an advantage over contextualism. A relativist can say that John believes the proposition that Picasso's *Guernica* is beautiful, and that Mary believes the negation of that proposition. On the other hand, it is unclear whether this amounts to a genuine disagreement given that the proposition that Picasso's *Guernica* is beautiful is true relative to John's standards and false relative to Mary's standards. Whether this counts as a disagreement depends on what happens to disagreement when propositional truth is relative in the relevant sense.

It is worth pointing out that this is sometimes put in terms of relativists being in a position to make sense of cases of so-called 'faultless disagreement'. A case of faultless disagreement is a disagreement in which neither party is making a mistake. Kölbel (2002, 2004) has emphasised this point, and he has argued that it is an important point in favour of relativism that it allows us to make sense of faultless disagreement. Even if we take Mary and John to disagree, there is perhaps a sense in which we want to say that neither of them is making a mistake. Relativists can make sense of that by claiming that their beliefs are true relative to their respective aesthetic standards. What John believes is true relative to his aesthetic standards, and what Mary believes is true relative to her aesthetic standards. However, the notion of 'faultless disagreement' is contentious.¹¹ It is also not clear that we need to make essential use of it in order to motivate relativism. For the most part, it will play a limited role in the following discussion.

There is a fairly natural view of disagreement which would seem to vindicate relativism. On this view, disagreement is simply a matter of one party believing a

¹⁰My preferred view would be to reject the second premise. See Chapter 2 for relevant discussion.

¹¹For critical discussion of faultless disagreement, see e.g. Stojanovic (2007) and Rosenkranz (2008).

proposition and another party believing its negation.

- (8) Two parties disagree if and only if there is a proposition such that one party believes that proposition and the other party believes its negation.

For our purposes, what matters is the claim that it is sufficient for two parties to disagree that there is a proposition such that one of them believes that proposition and the other believes its negation. If (8) gives us a sufficient condition for disagreement, relativists can say that Mary and John disagree in virtue of John believing the proposition that Picasso's *Guernica* is beautiful and Mary believing its negation. That is true even though that proposition is true relative to John's standards and false relative to Mary's standards. In other words, if relativists can accept something like (8), they are in a position to deliver the desired predictions as far as disagreement is concerned. The question is whether this combination of views is ultimately plausible. While it may be plausible that (8) gives us a sufficient condition for disagreement when propositional truth is not relative, one might start to doubt that this is the case when propositional truth is relative. In order to find an answer to questions like these, it makes sense to look at what happens to disagreement if propositional truth is relative to other parameters, such as times and worlds.

1.3 A Problem for the Simple Relativist Picture

A worry for the simple relativist picture is that the view of disagreement under discussion looks less plausible if we look at other views on which propositional truth is relative.¹² MacFarlane (2007) has argued that if there are tensed or temporally neutral propositions, propositions that are true or false relative to times, disagreement cannot simply be a matter of there being a proposition that one party believes while the other party believes its negation.¹³ Suppose that

¹²This argument plays a role in both Dreier's (2009) and Francén's (2010) discussion of relativist accounts of disagreement. However, for the most part I will be basing my presentation of these issues on MacFarlane's (2007) discussion.

¹³MacFarlane (2007) does not present the view he is arguing against as the view that it is sufficient for disagreement that there is a proposition such that one party believes that proposition and the other believes its negation. Instead he prefers to talk about there being a proposition that one party accepts and the others rejects. For the purpose of the present discussion, this does not

John at 2 p.m. believes the tensed proposition that Joe is sitting and that Mary at 3 p.m. believes the tensed proposition that Joe is not sitting. While Mary believes the negation of the proposition that John believes, this does not seem to be sufficient for them to disagree.¹⁴

MacFarlane argues that this means that a view of disagreement like (8) needs to be revised. He suggests that if we want to acknowledge the existence of tensed propositions, a natural replacement for (8) would be something like (9).¹⁵

- (9) Two parties disagree if and only if (i) there is a proposition such that one party believes that proposition and the other party believes its negation and (ii) the beliefs cannot both be accurate.

For our purposes, the precise details do not matter. The basic idea is that the reason Mary and John do not disagree, is that there is nothing that prevents both of them from being right. Even if John is right when he at 2 p.m. believes that Joe is sitting, this does not prevent Mary from being right when she at 3

make any difference.

¹⁴I am willing to grant that there might be different senses of 'disagree'. In that case we need to be careful when we ask whether Mary and John 'disagree'. However, as far as I can see, there is no sense of 'disagree' such that it is sufficient for the truth of 'Mary and John disagree' that John at 2 p.m. believes the tensed proposition that Joe is sitting and Mary at 3 p.m. believes the tensed proposition that Joe is not sitting.

¹⁵MacFarlane (2007) argues that (9) is not strictly speaking true as it stands. In fact, he argues that both the right-to-left and the left-to-right direction of (9) fail. Suppose that Mary at noon believes the proposition that the number of flies in the room is either odd or even, and that Tom at midnight believes the negation of that proposition. Their beliefs cannot both be accurate since Tom's belief cannot be accurate, he believes something that is necessarily false. However, MacFarlane claims that this is not sufficient for them to disagree. This means that (9) does not give us a sufficient condition for disagreement.

Furthermore, let us suppose that Mary at noon believes the tensed proposition that Socrates is sitting and that Peter at midnight believes the tensed proposition that Socrates was sitting twelve hours ago. In that case we want to say that they disagree, but according to (9) they do not. The reason is that there is not a proposition such that Mary believes that proposition and Peter believes its negation. This means that (9) does not give us a necessary condition for disagreement. However, for our purposes the important question is not what the replacement for (8) should look like, but whether it needs to be replaced. Whatever the right replacement for (8) turns out to be, it cannot be sufficient for two people to disagree that there is a proposition such that one of them believes that proposition and the other believes its negation, at least not when the beliefs concern different times.

p.m. believes that Joe is not sitting. All it takes is that Joe is sitting at 2 p.m. and that at 3 p.m. he is not sitting. We can follow MacFarlane and cash this out in terms of accuracy. A belief is accurate in the relevant sense if the proposition believed is true relative to the index that is relevant to assessment of the belief in its context.¹⁶ In the case of tensed propositions, the index will at least include the relevant time.

If we adopt (9) instead of (8), we can avoid the unfortunate prediction that Mary and John disagree when John at 2 p.m. believes the proposition that Joe is sitting and Mary at 3 p.m. believes the negation of that proposition. However, it is not clear that this goes well with the simple relativist treatment of (6) and (7).

(6) Picasso's *Guernica* is beautiful.

(7) Picasso's *Guernica* is not beautiful.

On the simple relativist story we can say that John believes the proposition that Picasso's *Guernica* is beautiful and that Mary believes the negation of that proposition. However, there is nothing that prevents their beliefs from both being accurate. Indeed, the point was that the proposition is true relative to John's standards, but false relative to Mary's standards. If we adopt (9) instead of (8), the simple relativist account does not have any discernible advantages over a more traditional contextualist account as far as disagreement is concerned.

1.3.1 Worlds Instead of Times

MacFarlane also argues that the same point can be made by relying on the less controversial assumption that propositions are true or false relative to worlds. Suppose that Jane, an inhabitant of the actual world, believes the proposition that Mars has two moons, and that June, who is Jane's counterpart and an inhabitant of another possible world, believes the negation of that proposition. Even if this means that there is a proposition such that Jane believes that proposition and June believes its negation, it still seems that they do not disagree. Again, if that is correct, it means that (8) cannot be right as it stands. Instead we may adopt

¹⁶Insofar as I am talking about indices rather than circumstances of evaluation, I am using Lewis' (1980) terminology, not Kaplan's (1989). For our purposes, we can treat these notions as equivalent.

something like (9), but then it is unclear whether the simple relativist account does any better than a contextualist account.

Since the assumption that propositions are true or false relative to worlds is less controversial than the assumption that propositions are true or false relative to times, it is worth spending some more time on this argument. It turns out to be a somewhat tricky question whether the argument requires any realist assumptions about modality in order to go through. Cappelen and Hawthorne (2009) have raised worries along these lines.¹⁷ They complain that, unless we accept a realist view of modality of the sort endorsed by Lewis (1986), the situation described is not a situation in which two individuals do not disagree even though one of them believes a proposition and the other believes its negation.¹⁸ They argue that one should not talk about an individual in another world as if it were an individual in a different country. If one says that there is a world in which an individual believes some proposition, this only means that it could be the case that there is an individual that believes the proposition. It does not mean that there is an individual that believes the proposition.

Even if it turns out that one must assume modal realism in order to get the argument going, this does not in and of itself show that the argument is not sound. However, since any realist assumptions about modality are likely to be controversial, this would make the argument less dialectically effective.

MacFarlane anticipates this sort of worry and argues that he can make the point without relying on any realist assumptions about modality. Instead of talking about Jane's counterpart in some other possible world, we need to consider whether Jane, in believing what she would have believed in a certain counterfactual situation, would disagree with what she actually believes. According to MacFarlane, it is not sufficient for her to disagree with what she actually believes that she would have believed the negation of the proposition she actually believes. However, as Cappelen and Hawthorne point out, this requires us to shift the focus away from disagreement between individuals and to start talking about

¹⁷Cappelen and Hawthorne (2009) talk about agreement rather than disagreement, but for our purposes it does not make any difference whether we are talking about agreement or disagreement.

¹⁸If one subscribes to presentism, the view that only temporally present objects exist, one might have similar concerns about what MacFarlane says about tensed propositions. See e.g. Cappelen and Hawthorne (forthcoming).

disagreement with what someone says or believes. Insofar as we are more interested in the former, that is not unproblematic. More importantly, Cappelen and Hawthorne argue that we should not be too quick to accept MacFarlane's verdicts about the relevant counterfactuals. If June had believed the proposition that it is not the case that Mars has two moons, there is at least some temptation to say that she would have disagreed with what she actually believes, namely the proposition that Mars has two moons. I am inclined to agree that it is not clear what we ought to say about the relevant counterfactuals, and this makes MacFarlane's argument somewhat shaky.

Another, and potentially more promising, way of trying to avoid having to talk about individuals inhabiting different worlds is to look for an individual who has a belief that is evaluated with respect to a world other than the actual world. The view of fiction explored by Predelli (2005, ch. 2) suggests that this may be possible. He asks us to consider an utterance of (10) that takes place as a part of a discussion of Milos Forman's film *Amadeus*.

(10) Salieri commissioned the *Requiem*.

Insofar as Salieri commissioned the *Requiem* in the film it would appear that (10) is true as uttered in a context like this. However, we may assume that in the actual world Salieri did not commission the *Requiem*. In that case, (10) would be false as uttered in a discussion about the actual events surrounding the *Requiem*. One way of respecting these judgements is to say when (10) is uttered as a part of a discussion of Milos Forman's film, it is not evaluated with respect to the actual world, but with respect to some other world, or more likely a set of worlds, determined by the fiction.¹⁹

Predelli focuses on language, but for our purposes it is what is going on at level of thought that matters. Suppose that John sincerely utters (10) as a part of a discussion about the film. In that case we may want to say that (10) expresses the proposition that Salieri commissioned the *Requiem* and that John believes that proposition. However, the world that is relevant for the assessment of his belief is not the actual world, but a world or set of worlds determined by the fiction. Let

¹⁹There will most likely be a certain measure of indeterminacy when it comes to the world, or worlds, determined by the fiction. However, it is not clear that this should be more problematic in this case than in other cases.

us further suppose that Mary sincerely utters (11) as a part of a discussion about who actually commissioned the *Requiem*.

(11) Salieri did not commission the *Requiem*.

In that case, we may want to say that (11) expresses the proposition that Salieri did not commission the *Requiem* and that Mary believes that proposition. But in this case, the world that is relevant for the assessment of her belief is simply the actual world. If all of this is correct, and that is admittedly assuming quite a lot, we would have a situation in which John believes the proposition that Salieri commissioned the *Requiem* and Mary believes its negation. Still, we do not want to be forced to say that Mary and John disagree.

The relevant view of fiction is obviously controversial, and there are a number of alternatives.²⁰ One might instead insist that (10) cannot be literally true, and that it only serves to pragmatically convey something true in the relevant context. Alternatively, one might take (10) to be elliptical for (12) with ‘It is true in the film *Amadeus*’ serving as a fiction operator.²¹

(12) It is true in the film *Amadeus* that Salieri commissioned the *Requiem*.

A full discussion of these issues is beyond the scope of present discussion. In any case, the point is not to present a conclusive argument for the conclusion that (8) cannot be true if propositions are true or false relative to worlds. There are interesting issues here that deserve to be explored further, but as it stands there does not appear to be a straightforward way of getting MacFarlane’s argument to work without some rather controversial assumptions. Again, this does not necessarily mean that the argument is not sound, but it does mean that it rests on assumptions that not everyone will want to accept. For the purpose of the following discussion, I will not assume that we need to reject (8) if propositions are true or false relative to worlds, but I will not assume that (8) is compatible with propositions being true or false relative to worlds either. Instead I will focus on views on which propositions are true or false relative to other parameters that are less likely to be associated with controversial metaphysical assumptions.

²⁰If one prefers to let propositions be specific with respect to the relevant world, one can say that (10) and (11) expresses different propositions about different worlds. See Schaffer (forthcoming-b) for a view on which propositions are world-specific in this way.

²¹See e.g. Lewis (1978).

1.3.2 Other Parameters

While MacFarlane (2007) focuses on worlds and times, the same point can be made if propositions are true or false relative to other parameters. The point emerges quite clearly if propositions are construed as sets of centred worlds. This view of propositions is typically associated with the picture of *de se* attitudes put forward by Lewis (1979).²² Centred worlds can be thought of as pairs that consist of a world and an individual.²³ A centred world $\langle x', w' \rangle$ is compatible with what x believes in w just in case it is compatible with what x believes in w that she is x' in w' . One believes a proposition, a set of centred worlds, just in case every centred world compatible with what one believes is a member of the set.

On this view, propositions can have different truth-values relative to different individuals. For instance, the proposition that is the set of centred worlds in which the centre is hungry, may be true relative to John, but false relative to Mary. However, if John believes this proposition and Mary believes its negation, that is not sufficient for them to disagree. To use Lewis' terminology of self-ascription, this would essentially be situation in which John self-ascribes the property of being hungry and Mary self-ascribes the property of not being hungry. That is not enough for them to disagree. This is particularly interesting insofar as this picture has provided inspiration for contemporary relativists.²⁴ While this does not show that it is wrong for relativists to draw inspiration from such views, one should be careful not to overstate the similarities between the views.

The point can also be clearly seen if there are propositions that have different truth-values relative to different locations. One might claim that a sentence like (13) expresses a location-neutral proposition, a proposition that is true relative to some locations, but false relative to others.²⁵

²²Lewis (1979) primarily talks about the objects of mental states as properties, but for our purposes this does not make any difference whether we talk about the objects of mental states as properties or sets of centred worlds. For the purpose of presenting the view, I am for the most part following Ninan (2010).

²³Centred worlds are typically thought of as triples that consist of a world, an individual, and a time. Since we have already discussed views on which propositions are true or false relative to times, I will ignore the time coordinate for the purpose of the following discussion.

²⁴See e.g. Egan (2007, 2010) for a relativist account who is inspired by Lewis' (1979) work on *de se* attitudes.

²⁵See e.g. Recanati (2007) for relevant discussion of location-neutral propositions. See e.g. also

(13) It's raining.

Suppose that John, who is in St Andrews, believes the location-neutral proposition expressed by (13) and that Mary, who is in Oslo, believes the negation of that proposition. Still, that does not seem to be sufficient for them to disagree. After all, there is nothing that prevents Mary and John from both being right if it is raining in St Andrews and not in Oslo.

Focusing on parameters like this allows us to avoid some of the metaphysical issues involving modality. On the other hand, the view that there are propositions that have different truth-values relative to different times, individuals, or locations, is more controversial than the view that there are propositions that have different truth-values relative to different worlds. There is considerable opposition to the idea that propositions can be time or location-neutral in this sense, and some philosophers doubt that we can make sense of the idea that propositions are true or false relative to different times or locations. Frege (1918) is a classic example of a philosopher who denies that there are propositions that vary in truth-value across times.

But are there not thoughts which are true today but false in six months' time? The thought, for example, that the tree is covered with green leaves, will surely be false in six months' time. No, for it is not the same thought at all. The words 'This tree is covered with green leaves' are not sufficient by themselves to constitute the expression of a thought, for the time of utterance is involved as well. Without the time-specification thus given we have not a complete thought, i.e. we have no thought at all. Only a thought with the time-specification filled out, a sentence complete in every respect, expresses a thought. But this thought, if it is true, is true not only today or tomorrow but timelessly. (Frege, 1918, p. 343)

A more recent example of someone who does not think that we can make sense of proposition that vary in truth-value across times is (King, 2003). He expresses this sentiment quite clearly.

Kaplan (1989).

[...] though it seems correct to hold that the things I believe, doubt, etc. can change truth value across worlds (i.e. some of the things I believe are true though they would have been false had the world been different), it is hard to make sense of the idea that the things I believe may change truth value across time and location. What would it be e.g. to believe that the sun is shining, where what I believe is something that varies in truth value across times and locations in the actual world? It seems clear that when I believe that the sun is shining I believe something about a particular time and location, so that what I believe precisely does not vary in truth value over times and locations. (King, 2003, p. 196)

This is not the time or place to attempt to adjudicate these disputes. It is sufficient to note that there is considerable controversy surrounding these issues. Having said that, one would not expect a relativist to say that it does not make sense for the things we believe to have different truth-values relative to different times or locations. After all, one might have similar doubts about whether it makes sense for the things we believe to be true or false relative to different aesthetic standards.

Where does all of this leave us with respect to the simply relativist story? We can state the general objection as follows: When we look at other views on which propositions are true or false relative to some parameter, disagreement does not work the way that relativists want it to work in the case of (6) and (7).

(6) Picasso's *Guernica* is beautiful.

(7) Picasso's *Guernica* is not beautiful.

In particular, it does not seem to be sufficient for two parties to disagree that there is a proposition such that one party believes that proposition and the other believes its negation. Why should we then expect disagreement to work the way that relativists want it to work in the case of (6) and (7)? Furthermore, if we replace (8) with something like (9) across the board, it appears that the simple relativist treatment of (6) and (7) no longer offers any advantages over a contextualist treatment when it comes to dealing with disagreement. I take it that this is a conclusion that relativists need to resist. In what follows I will discuss ways in which relativists can respond to this argument.

1.4 Two Relativist Replies

At this point I will explore two lines of response on behalf of the relativists. The first option is to argue that propositions are not true or false relative to parameters like worlds, times, individuals, and locations. The second option is to argue that the parameters that relativists are talking about are different from parameters like world, times, individuals, and locations. MacFarlane (2007) can be thought of as advocating something like the second response. However, I will argue that the first option should also be taken seriously.

1.4.1 The First Reply

While it might not be very sensible for relativists to claim that it does not make sense for the things we believe to have different truth-values relative to different times or locations, that does not mean that they have to think that propositions are in fact true or false relative to such parameters. They could deny that propositions are true or false relative to times and locations, and continue to argue that they are true or false relative to, say, aesthetic standards. As was noted earlier, it is not clear whether this would also have to apply to worlds as well as times and locations. However, even if it also applies to worlds, there does not appear to be any special commitments associated with relativism that prevents relativists from thinking of propositions as world-specific in the relevant sense. On this kind of view, a proposition would be about a particular world in more or less the same way that King (2003, p. 196) argues that it is about a particular time and location. This would allow relativists to hold on to (8) and the simple relativist treatment of sentences like (6).

(6) Picasso's *Guernica* is beautiful.

One might worry that this somehow goes against the spirit of relativism. To a certain extent relativists can be seen as building on views like Lewis' (1979) view of *de se* attitudes and work in formal semantics by, among others, Lewis (1980) and Kaplan (1989). However, even if relativists are inspired by views on which truth is relative to, say, worlds and times, it is not clear why relativists need to be committed to truth being relative to other parameters when the motivation for positing these parameters is different from the parameters that relativists want to

posit. The lesson might be that we should not overstate the similarities between the kind of relativists views that we are discussing here and views that look similar from a formal point of view.

Furthermore, it may not be necessary to say that there is no sense of ‘proposition’ such that there are propositions that are true or false relative to different worlds, times, individuals, or locations. Instead one might adopt a view on which there are multiple levels of content with different kinds of content serving different purposes.²⁶ This would allow a relativist to leave room for something like Lewis’ (1979) view of *de se* attitudes. The level of content that relativists are interested in, is the level of content that is relevant for assessing whether people agree or disagree. In this sense of ‘proposition’, propositions are not true or false relative to times and locations, and there is no reason to give up (8).

This is not to say that this line of response is unproblematic. One way or another, this line of response comes with significant commitments. Either the relativist will be committed to propositions not being true or false relative to times and locations, or she will be committed to there being multiple levels of content. Either way, a lot more work remains to be done.

Even if this is a plausible response to the argument that was presented in the previous section, one might have other reasons for doubting that the relativist picture of disagreement is on the right track. For instance, when one tries to motivate the judgement that Mary and John do not disagree when John at 2 p.m. believes the tensed proposition that Joe is sitting and Mary at 3 p.m. believes its negation, one is likely appeal to the fact that Mary and John could both be right. This is not a motivation that sits well with the simple relativist story. After all, that could just as well be said about Mary and John when they have different aesthetic standards. What John believes is true relative to his standards, and what Mary believes is true relative to her standards.

However, it is doubtful whether any clear-cut argument against the relativist picture of disagreement is likely to emerge from considerations like this. The claim that Mary and John cannot disagree unless one of them is wrong, is likely to be contested. As noted earlier, Kölbel (2002, 2004) has argued that an important motivation of relativism is that it allows us to make sense of cases of so-called

²⁶For a defence of a view on which there are multiple levels of content, see e.g. Recanati (2007).

‘faultless disagreement’, cases of disagreement in which neither party is wrong. Many have expressed doubts about whether there could be genuine cases of faultless disagreement, but it is a different matter to provide an argument against faultless disagreement that does not just assume that it is a necessary condition for disagreement that the parties cannot both be right.²⁷

While there might not be a decisive argument against the relativist picture of disagreement, a further question is whether it is possible to offer any independent motivation for this way of thinking about disagreement. What, if anything, can be said in favour of thinking that Mary and John disagree when he believes the proposition that Picasso’s *Guernica* is beautiful and she believes its negation? While (8) may provide a plausible sufficient condition for disagreement when propositional truth is not relative, it is not clear that this is the case if relativism is true. If propositions are not true or false relative to any other parameters, such as times or locations, it is hard to find counterexamples to the picture of disagreement that goes with the simple relativist story, but it is also hard to find independent evidence for it. After all, where is that evidence going to come from? The only other cases of disagreement involve propositions that do not vary in truth-value across any parameters.

One way of trying to motivate the relativist picture of disagreement, is to argue that disagreement is tied to attributions of truth and falsity. According to the relativists, the proposition that John believes is false relative to the aesthetic standards that are in play in Mary’s context. They predict that (14) would be true as uttered by Mary.

(14) What John believes is false.

On this picture it makes sense for Mary to believe that what John believes is false. Furthermore, as Cappelen and Hawthorne (forthcoming) observe, it is really strange to say things like (15).

(15) Mary believes that what John believes is false, but she does not disagree with him.

²⁷See e.g. Rosenkranz (2008) for an attempt to construct an argument against faultless disagreement.

This could be taken as evidence that disagreement is tied to attributions of truth and falsity in such a way that Mary and John can be expected to disagree even if relativism is true. If there was no such connection between disagreement and attributions of truth and falsity, there would be no obvious reason to expect (15) to be an inappropriate thing to say. A principle like (16) would explain why (15) is such a strange thing to say and could be used to motivate the claim that Mary and John disagree.

- (16) For every x and every y , if x believes that what y believes is false, x disagrees with y .

If there were propositions that are true or false relative to, say, times or locations, (16) would be problematic. Insofar as we are already assuming that propositions are not true or false relative to such parameters, that may not seem to be a problem. From this point of view, I am willing to grant that this provides some independent evidence for the claim that Mary and John disagree. However, if a relativist wants to buy into the idea that there are multiple levels of content, she would have to deny that expressions like ‘what she believes’ can pick out the contents that have different truth-values relative to different times and locations.²⁸

Let us take stock. One way for relativist to respond to the argument that was presented in §1.3, is to deny that there are propositions that are true or false relative to parameters like times and locations. This would allow relativists to hold on to a simple view of disagreement. However, there are serious commitments associated with this line of response. Relativists will either have to argue that propositions are not true or false relative to parameters such as times and locations, or to argue that there are multiple levels of content.

While this provides a response to the argument that was presented in §1.3, it is not clear how one can provide independent motivation for the relevant view of disagreement. A promising way of doing that may be to appeal to the con-

²⁸In any case, we should not exaggerate the importance of considerations like this. While (16) may be plausible if propositional truth is not relative, one might have doubts about whether this is still the case if relativism is true. After all, it is plausible that (8) provides a sufficient condition for disagreement when propositional truth is not relative, but that is not obvious if relativism is true. Since propositional truth is only relative to the parameters that relativists want to posit, the only other cases that we can take into consideration are cases in which the relevant propositions do not vary in truth-value across any parameters.

nection between disagreement and attributions of truth and falsity, but insofar as we are discussing views on which truth is relative, such a connection may be contentious. It might be that we do not have much to go on when it comes to providing independent evidence for or against the relevant view of disagreement.

1.4.2 The Second Reply

There might be all sorts of reasons why a relativist would not want to pursue the first reply. For instance, she might be committed to time and location-neutral propositions and unwilling to countenance multiple levels of content. In that case, another option is to argue that the different parameters that truth is relative to, work differently with regard to disagreement. They may grant that it is not sufficient for Mary and John to disagree that he believes the location-neutral proposition that it's raining and she believes its negation. However, even if the proposition that Picasso's *Guernica* is true relative to John's standards and false relative to Mary's standards, it may still be sufficient for Mary and John to disagree that John believes that proposition and Mary believes its negation.

As it stands, this response sounds blatantly *ad hoc*. Why should we think that disagreement works differently when the relevant propositions have different truth-values relative to different aesthetic standards rather than, say, times or locations? However, a relativist does not have to leave it at that. To make the move seem less *ad hoc* they can attempt to provide reasons for thinking that the relevant parameters should be treated differently. MacFarlane (2005b, 2007) has developed a relativist framework that appears promising in this regard, and for the purpose of the present discussion, I will focus on his account.

A central part of MacFarlane's account is the distinction between the context of use and the context of assessment. Roughly speaking, the former is a context in which a sentence is used and the latter is a context in which the use of a sentence is assessed. MacFarlane proposes that when we talk about sentences being true or false, we need a doubly contextual truth-predicate. We need to talk about a sentence being true, not just at a context of use, but at a context of use and a context of assessment. According to MacFarlane (2005b), for relativism to be true is for there to be a sentence that has different truth-values relative to different contexts of assessment. For instance, (6) could be true relative to a

context of assessment in which John is the assessor, but false relative to a context of assessment in which Mary is the assessor.

(6) Picasso's *Guernica* is beautiful.

We can distinguish between two different roles that the context of use or the context of assessment can play in determining the truth-value of a sentence. While the truth-value of a sentence can depend on the context insofar as the context plays a role in determining the proposition expressed by the sentence, it can also depend on the context insofar as it plays a role in determining the value of a parameter that the relevant proposition is true or false relative to.

Since we are focusing on propositional truth, it is the latter that is relevant for our purposes. MacFarlane's framework allows us to distinguish between parameters that are initialised by the context of use and parameters that are initialised by the context of assessment. According to a relativist view of this sort, propositions are true or false relative to a parameter that is initialised by the context of assessment rather than the context of use. In other words, if a relativist treatment of a sentence like (6) is correct, it expresses a proposition that is true or false relative to the aesthetic standards determined by the context of assessment.

Since some parameters are initialised by the context of use and some are initialised by the context of assessment, there is perhaps a basis for saying that disagreement works differently when different parameters are involved. In order to provide an account according to which disagreement works differently when different parameters are involved, MacFarlane (2007) takes it to be an assessment-sensitive matter whether a belief is accurate. To see how this works, let us suppose that propositions are true or false relative to worlds and times, in addition to aesthetic standards. The world and time are determined by the context of use, and the aesthetic standards are determined by the context of assessment. In that case, a belief is accurate relative to a context of assessment just in case the proposition believed is true relative to the world and time of the context of use and the aesthetic standards of the context of assessment. If we adopt something like (17), this would allow the relativists to make the desired predictions.

(17) Two parties disagree relative to a context of assessment if and only if (i) there is a proposition such that one party believes that proposition and

the other party believes its negation and (ii) the beliefs cannot both be accurate relative to that context of assessment.²⁹

According to this view of disagreement, it matters whether the relevant parameters are initialised by the context of use or by the context of assessment. If John at 2 p.m. believes the tensed proposition that Joe is sitting and Mary at 3 p.m. believes the negation of that proposition, their beliefs can both be accurate. It does not matter what the context of assessment is like, since the relevant time is determined by the context of use. However, if John believes the proposition that Picasso's *Guernica* is beautiful and Mary believes its negation, there is no context of assessment such that their beliefs are both accurate. The proposition and its negation cannot both be true relative to the same aesthetic standards.

If this is correct, there is some basis for treating the parameters that relativists want to posit differently than parameters like worlds and times. This means that the relativists have a response to the argument that was presented in §1.3. Even if disagreement works in a certain way when different worlds, times, and locations are involved, that does not entail that it has to work in the same way when different aesthetic standards are involved.

Again, it is a further question whether there is independent evidence for this view of disagreement. In particular, are there any independent reasons for thinking that (17) is the right view of disagreement when there are propositions that are true or false relative to parameters that are initialised by the context of assessment? One way of putting the question is to ask why we should adopt (17) instead of (18).

- (18) Two parties disagree if and only if (i) there is a proposition such that one party believes that proposition and the other party believes its negation and (ii) the beliefs cannot both be accurate relative to the parties' respective contexts of assessment.

What this means is that if Mary and John disagree, it cannot be the case that John's belief is accurate relative to his context of assessment and that Mary's belief is accurate relative to her context of assessment. That requirement is not satisfied in

²⁹Strictly speaking, this means that disagreement is assessment-sensitive. However, this turns out to be irrelevant. As MacFarlane (2007) points out, if two parties disagree relative to a context of assessment, they disagree relative to every context of assessment.

a situation in which John believes the proposition that Picasso's *Guernica* is beautiful and Mary believes its negation if that proposition is true relative to John's context of assessment, but false relative to Mary's context of assessment. In that case, John's belief would be accurate relative to his context of assessment and Mary's belief would be accurate relative to her context of assessment. From the point of view of (18), it does not matter whether the parameters are initialised by the context of assessment or the context of use.

Perhaps it is too much to ask that MacFarlane should provide a reason for adopting (17) instead of (18). MacFarlane offers a view of disagreement that allows us to make sense of apparent cases of disagreement that we may otherwise find it difficult to make sense of, and perhaps we should not expect more. Furthermore, even if there is no independent evidence that supports (17), it is not clear that there is any evidence against it either. Having said that, it would certainly be interesting if there were any reasons for adopting (17) instead of (18).

Could a relativist try to appeal to the alleged connection between disagreement and attributions of truth and falsity? If we are talking about propositional truth, there is a worry that this does not allow us to distinguish between the parameters that relativists want to posit and parameters like worlds, times, and locations. Suppose that John, who is in St Andrews, believes the location-neutral proposition that it's raining, and that Mary, who is in Oslo, at believes its negation. Furthermore, let us suppose that it is raining in St Andrews, but not in Oslo. In that case, (14) would be true at as uttered by Mary in Oslo.

(14) What John believes is false.

What John believes is the proposition that it's raining, and that proposition is false relative to Oslo. However, even if Mary believes that what John believes is false, we do not want to conclude that they disagree. As far these attributions of propositional truth are concerned, there is no difference between propositions that are true or false relative to contexts of use and propositions that are true or false relative to contexts of assessment. It is true that Mary can say that what John believes is false when he believes the proposition that Picasso's *Guernica* is beautiful, and that is false relative to her standards. However, that is true regardless of whether the aesthetic standards are determined by the context of assessment

or the context of use. Furthermore, it cannot be sufficient for Mary and John to disagree unless one is prepared to accept that this is also true in cases involving time or location-neutral propositions.³⁰ MacFarlane (2007) recognises that he needs to give up (16) if he wants to countenance the existence of propositions that are true or false relative to parameters like times or locations.

- (16) For every x and every y , if x believes that what y believes is false, x disagrees with y .

If a relativist finds that objectionable, she may take that to be a reason to go with the first reply and deny that there are propositions that are true or false relative to times or locations. In any case, it makes it difficult to use attributions of propositional truth as a basis for arguing that Mary and John disagree when different aesthetic standards are involved.

Perhaps it is more promising to look for a connection between disagreement and attributions of truth as it applies to utterances, as opposed to propositions.³¹ Following MacFarlane (2009), we can adopt (19) as a definition of utterance truth.³²

- (19) An utterance of a sentence s in a context of use c_U is true relative to a context of assessment c_A if and only if the proposition expressed by s in c_U is true relative to the world, time, and location of c_U and the aesthetic standards of c_A .

In this case it makes a difference whether the parameters are initialised by the context of assessment or the context of use. Suppose that (13) expresses a propo-

³⁰As a result of complications involving tense, matters are a bit more complicated in the case of tensed propositions. See e.g. Cappelen and Hawthorne (forthcoming) for relevant discussion.

³¹From the point of view of thinking about disagreement, it would perhaps make sense to talk about truth as it applies to belief-states. However, I find it easier to talk about utterance truth, and I will continue to do so for the purpose of this discussion. As far as I can see, nothing substantial turns on this.

³²MacFarlane (2005b, 2007, 2009) does not seem to be too happy about applying truth to utterances. He takes the notion of utterance truth to be a technical notion. His notion of 'accuracy' can play a similar role. However, given the way that he uses it, it is not clear that it is any less of a technical notion than utterance truth. Furthermore, I am not entirely convinced that utterance truth is a purely technical notion, in the sense that we have no independent grip on it. I will therefore continue to talk about utterance truth.

sition that has different truth-values relative to different locations and that John utters that sentence in St Andrews.

(13) It's raining.

If Mary is in a location in which it is not raining, she can say what John said is false. However, according to the definition of utterance truth we are working with, (20) would still be false as uttered by Mary.³³

(20) John's utterance of 'It's raining' is false.

However, now we can see a difference between locations and aesthetic standards. If (6) expresses a proposition that is false relative to Mary's aesthetic standards and John utters that sentence, Mary can say that his utterance is false. In that case, (21) would be true as uttered by Mary.

(21) John's utterance of 'Picasso's *Guernica* is beautiful' is false.

Does this give us a reason to think that Mary disagrees with John? While I do not want to deny that this is some evidence that they disagree, one should be careful about relying too much on considerations like this. After all, why should we expect disagreement to be tied to attributions of utterance truth when it is not tied to attributions of propositional truth? Even if there are no clear counterexamples to the claim that disagreement is tied to attributions of utterance truth, we have already acknowledged that our judgements about the connection between disagreement and attributions of propositional truth might be mistaken. That makes it somewhat questionable to rely on judgements about the connection between disagreement and attributions of utterance truth.

This suggests that it would be preferable to have an account that preserves the connection between disagreement, propositional truth, and utterance truth. That is possible if one denies that propositions are true or false relative parameters that are determined by the context of use. The only relevant parameters are parameters that are determined by the context of assessment. In effect, this would be an account that combines the first and the second reply.

³³If you think that there is something strange about propositional truth and utterance truth coming apart in this way, you are not alone. See e.g. Cappelen and Hawthorne (2009) and MacFarlane (2009) for relevant discussion. See also Chapter 5.

Let us take stock. One way for relativists to respond to the argument that was presented in §1.3 is to argue that disagreement works differently when different parameters are involved. While this may seem *ad hoc*, MacFarlane (2005b, 2007) has developed a relativist framework that allows us to distinguish between parameters that are initialised by the context of assessment and parameters that are initialised by the context of use. This provides a basis for treating different parameters differently when it comes to disagreement.

However, the extent to which the relevant view of disagreement can be said to be independently motivated is still not clear. While it is difficult to appeal to a connection between disagreement and attributions of propositional truth, it might be possible to appeal to a connection between disagreement and attributions of utterance truth. However, it is not clear that it is a good idea to invest too much in the connection between disagreement and attributions of utterance truth, when it has already been acknowledged that the connection between attributions of propositional truth and disagreement fails.

1.5 Indexical Relativism and Nonindexical Contextualism

So far I have focused on views on which propositional truth is relative. However, MacFarlane's (2005b) distinction between contexts of assessment and contexts of use allows us to identify another theoretical option. According to MacFarlane's preferred version of relativism, there are sentences that express propositions that have different truth-values relative to different contexts of assessment. An alternative to this view, is the view that the relevant sentences express different propositions relative to different contexts of assessment. Weatherson (2009) calls this view 'indexical relativism', and he argues that it gives us a good treatment of indicative conditionals.³⁴ Cappelen (2008a, 2008b) has explored a similar view that he calls 'content relativism'. It involves treating sentences as expressing different propositions relative to different contexts of interpretation. However, for the purpose of the following discussion I will be sticking with Weatherson's terminology and I will be glossing over any differences between his view and Cappelen's view.

³⁴MacFarlane (2005b) calls the view 'expressive relativism', but I will not be using his terminology here.

While indexical relativism may have lots of interesting applications, what matters for our purposes is how it deals with disagreement. The basic idea is that a sentence like (6) expresses one proposition relative to John's context of assessment and a different proposition relative to Mary's context of assessment.

(6) Picasso's *Guernica* is beautiful.

This should not be confused with the standard contextualist view that treats (6) as expressing different propositions in different contexts of use. Even if we hold the context of use fixed, (6) would still express different propositions relative to different contexts of assessment. Suppose that (6) is uttered by John. The proposition expressed relative to John's context of assessment is the proposition that Picasso's *Guernica* is beautiful relative to John's standards. However, the proposition expressed relative to Mary's context of assessment is the proposition that Picasso's *Guernica* is beautiful relative to Mary's standards. We may suppose that the proposition expressed relative to John's standards is true, while the proposition expressed relative to Mary's standards is false.

Does this provide a basis for claiming that Mary and John disagree? If we think about disagreement as a phenomenon at the level of thought, that is not clear. The most natural way to characterise their beliefs is to think of John as believing the proposition that Picasso's *Guernica* is beautiful relative to John's standards, and to think of Mary as believing the proposition that Picasso's *Guernica* is beautiful relative to Mary's standards. If that is correct, indexical relativism is in the same position as contextualism with respect to what is going on at the level of belief. This suggests that indexical relativism is no better off than contextualism when it comes to dealing with disagreement in this sense.

Having said that, indexical relativism appears to be better suited for dealing with other kinds of data that have played an important role in the debate between contextualists and relativists. For instance, if indexical relativism is correct, it makes sense for Mary to disagree with John's utterance and to reject his assertion. This could explain what is going in dialogues like (22).

- (22) a. John: Picasso's *Guernica* is beautiful.
b. Mary: I disagree./No, it is not beautiful.

If John's utterance of (22a) expresses a false proposition relative to Mary's con-

text of assessment, it would make sense for her to reject his assertion.

These kinds of dialogues are a useful way of eliciting judgements about disagreement. But it is also possible to ask whether Mary and John disagree when they are in different contexts and not speaking to each other. Insofar as this is a matter of what Mary and John believe, this is more complicated from the point of view of indexical relativism.³⁵

It is also worth noting that MacFarlane's way of defending relativism leaves room for a view he calls 'nonindexical contextualism'. Relativists and nonindexical contextualists agree that propositional truth is relative. The difference comes down to whether the relevant parameters are initialised by the context of assessment or the context of use. According to relativism, the relevant parameters are initialised by the context of assessment. According to nonindexical contextualism, the relevant parameters are initialised by the context of use.

MacFarlane (2009) does not take nonindexical contextualism to offer any advantages over contextualism when it comes to dealing with disagreement. The point of introducing contexts of assessment was to distinguish the parameters that relativists want to posit from parameters like worlds, times, and locations. If the aesthetic standard parameter was initialised by the context of use, it would be no different from the world or the time parameter. It would not be sufficient for Mary and John to disagree that John believes the proposition that Picasso's *Guernica* is beautiful and Mary believes the negation of that proposition.³⁶ Insofar as I am mainly interested in questions about disagreement, this makes nonindexical contextualism less interesting as an independent theoretical option. Still, from the point of view of providing a taxonomy of the relevant views, it is a position that should be recognised. If we accept MacFarlane's way of construing the debate, we are left with the following four positions:

- Contextualism

³⁵When Weatherson (2009) attempts to motivate an indexical relativist treatment of indicative conditionals using agreement data, he talks about speakers agreeing with an utterance. Indexical relativism delivers the right predictions with respect to this kind of agreement data, but there might be other kinds of agreement data that the view does not explain.

³⁶MacFarlane (2009) suggests that nonindexical contextualism has other advantages over contextualism, such as being in a position to provide a better account of belief and indirect speech reports. See e.g. also Brogaard (2008).

- Relativism
- Indexical Relativism
- Nonindexical Contextualism

However, insofar as the focus is on considerations involving disagreement, I will mainly be talking about contextualism and relativism. For the most part I will be ignoring indexical relativism and nonindexical contextualism as independent options.

Conclusion

The purpose of this chapter has been to introduce the issues involving relativism and disagreement. Beyond a commitment to propositions being true or false to a ‘non-standard’ parameter, I have not attempted to offer a general and uniform characterisation of relativism. What the relativist views I that am interested in have in common, is that they are motivated by considerations involving disagreement. In particular, relativism is supposed to do better than contextualism when it comes to dealing with disagreement.

I have discussed an objection to relativist accounts of disagreement that is based on how disagreement works when truth is relative to parameters like worlds and times. The upshot of the discussion is that there are at least two ways for relativists to respond to this argument. The first is to deny that propositions are true or false relative to parameters like worlds and times. The second is to say that the parameters that relativists want to posit are somehow different from the more traditional parameters. Both of these options seem to be viable, but in either case there is a further question about whether the view can be independently motivated. I have suggested that it might be possible for relativists to appeal to a connection between disagreement and truth in order to provide independent motivation for their account of disagreement, but such a connection might be problematic when truth is relative.

Perhaps there is a way to argue that the relativist account of disagreement is incorrect without relying on premises that the relativists are already committed to denying. Or perhaps it is possible to show that the relativist account of disagreement is independently motivated. At this point I am not going to make any

assumptions either way. In the following chapters I will focus on more general questions about agreement and disagreement, and the argument from disagreement as a problem for contextualism.

Chapter 2

Agreement and Disagreement

SOCRATES: What are the subjects of difference that cause hatred and anger?

(Plato, *Euthyphro* 7b)

Introduction

Discussions that concern agreement or disagreement on some level or another are common in philosophy. For instance, there has been a lot of recent interest among philosophers of language and linguists in the role that considerations involving agreement and disagreement play in motivating a relativist semantics for a wide range of expressions that includes epistemic modals, knowledge ascriptions and predicates of taste.¹ Among other things it has been argued that there are cases of faultless disagreement, cases of disagreement in which neither party is making a mistake, and that relativists are in a position to secure such faultless disagreement.

While considerations involving agreement and disagreement have played a particularly central role in these debates, there are lots of other debates in philosophy in which questions about agreement and disagreement are relevant. Moral philosophers discuss the alleged intractability of moral disagreement and

¹See e.g. Kölbel (2002, 2004), Richard (2004, 2008), Egan et al. (2005), Lasersohn (2005), Egan (2007, 2010), Stephenson (2007) and MacFarlane (2005a, 2005b, 2007, forthcoming-a, forthcoming-b).

whether certain theories about morality and moral thought and talk are incompatible or in tension with there being genuine moral disagreement. Epistemologists argue about the epistemic significance of disagreement and metametaphysicians debate whether disagreements about ontology are merely verbal. Extending the list is probably more a matter of patience than anything else.

In light of all of this it seems that we could benefit from a better understanding of agreement and disagreement. I am not pretending that all of the philosophical questions mentioned above will be answered if we succeed in achieving a better of understanding of these phenomena, but perhaps some headway can be made in certain areas. Therefore, what I want to do is to step back from the philosophical debates in which considerations involving agreement and disagreement play a role, and talk about agreement and disagreement in their own right. The central question that I want to address concerns the attitudes that are involved when we agree or disagree. Is it the case that agreement and disagreement always involve a particular kind of attitude or can it involve different kinds of attitudes? For instance, can all cases of agreement and disagreement be understood in terms of what the parties believe, or are there cases of agreement or disagreement that involve other attitudes?

My main goal is to attempt to answer these question by developing a view of agreement and disagreement that treats disagreement as involving conflicting attitudes and agreement as involving converging attitudes. I will be taking Stevenson's (1937, 1944, 1963) notion of 'conflicting attitudes' as my starting point. According to the view that I want to defend, agreement and disagreement is not exclusively a matter of what the parties believe. Agreement and disagreement can involve a wide range of attitudes such as wanting, liking, approving, and so forth. While this may sound like an ambitious project, I want to stress that it does not have to take the form of a fully reductive theory or analysis in terms of necessary and sufficient conditions. It is still possible to say something interesting about agreement and disagreement without engaging in such an overly ambitious project.

As a part of this overall project I will be arguing against attempts to understand all cases of agreement and disagreement in terms of what the parties believe. While this might appear to be an initially attractive view of agreement and disagreement, I want to argue that there are cases of agreement and disagree-

ment that are hard to reconcile with this view. Attempts to understand these cases in terms of what the parties believe run into problems. There might be ways of avoiding these problems while maintaining that agreement and disagreement is always a matter of what the parties believe. However, in order to do that it is necessary to take on board certain further commitments that are not unproblematic in their own right.

In §2.1 I distinguish between two questions concerning agreement and disagreement, the attitude question and the object question, and I make it clear that I am mostly interested in the former question. In §2.2 I discuss an answer to the attitude question that is based on the view that agreement and disagreement is always a matter of what the parties believe. In §2.3 and §2.4 I argue that there are cases of disagreement that should not be understood in terms of what the parties believe. In §2.5 I discuss the possibility that someone who adopts a relativist view may be able to avoid these problems. In §2.6 I propose that we think about agreement and disagreement in terms of converging and conflicting attitudes, and that this forms the basis for a better view of agreement and disagreement. In §2.7 I say something about what it takes for two parties to have conflicting attitudes. In §2.8 I discuss the implications for the debate about faultless disagreement.

2.1 The Questions

Some cases of agreement and disagreement are relatively straightforward. Let us start with a fairly simple and uncontroversial case of disagreement. Suppose that Harry, Mary, and John are having a discussion about 19th century Russian literature. Harry wants to know who wrote *War and Peace* and the following exchange takes place:

- (23) Harry: Who wrote *War and Peace*?
- a. Mary: *War and Peace* was written by Tolstoy.
 - b. John: No, it was written by Dostoyevsky, not Tolstoy.

Let us assume that the parties are sincere and competent users of the language who are not confused about what they believe in any relevant sense. In that case, it seems clear that Mary and John disagree in this example. But what is going on here? What can we say about the disagreement between Mary and John? An

initially reasonable reaction is to say that there is a proposition, the proposition that Tolstoy wrote *War and Peace*, that Mary believes to be true and John believes to be false. This may or may not be the whole story, but it goes at least some way towards explaining why Mary and John disagree.

Let us also look at a similarly straightforward case of agreement. Suppose that Harry, Mary, and John are talking about chess. They start talking about the former world chess champion Garry Kasparov, and Harry wants to know when he first became world champion. He decides to ask Mary and John, and the following exchange takes place:

- (24) Harry: When did Kasparov become world chess champion?
 - a. Mary: Kasparov became world chess champion in 1985.
 - b. John: Yes, he became world chess champion in 1985.

I take it that Mary and John agree in this example. Moreover, something similar to what was said about Mary and John's disagreement in (23) can be said about Mary and John's agreement in (24). In this case, there is a proposition, the proposition that Kasparov became world chess champion in 1985, that both Mary and John believe to be true. Cases like (23) and (24) provide a useful point of reference for the following discussion insofar as they provide us with relatively simple and straightforward cases of agreement and disagreement.

The central question that I want to address concerns the attitudes that are involved when we agree or disagree. For instance, it is natural to think of Mary and John's disagreement in (23) and Mary and John's agreement in (24) as having something to do with what the parties believe. In the case of Mary and John's disagreement in (23), Mary believes the proposition that *War and Peace* was written by Dostoyevsky while John believes the negation of that proposition. Similarly, in the case of Mary and John's agreement in (24) they both believe the proposition that Kasparov became world chess champion in 1985.

The question is whether there is anything general and interesting to be said about the attitudes that we have when we agree or disagree. Is it the case that agreement and disagreement always involve belief, as in the case of (23) and (24), or can there be cases of agreement and disagreement involving other attitudes? Let us call this the 'attitude question'.

The Attitude Question: What kinds of attitudes may be involved when two parties agree or disagree?

Behind this way of stating the question there is a presupposition that whether two parties agree or disagree is at least partly a matter of which attitudes the parties have. This presupposition may of course be challenged, but it seems to be a reasonable assumption. There is certainly something plausible about the thought that whether we agree or disagree somehow depends on our attitudes. When we want to know whether two parties agree or disagree, we need to know what their attitudes are like. This is not to say there cannot be other relevant factors, but the attitudes of the parties are always relevant.

It is important to make it clear that this means that I am thinking about agreement and disagreement primarily as phenomena at the level of mental states, not at the level of language. When we ask whether two individuals agree or disagree, this is a matter of which attitudes they have towards the relevant objects, not a matter of what they are saying. This point of view is endorsed by Jackson and Pettit (1998). They put the point quite succinctly: ‘Moral disagreement, and indeed disagreement in general, is a psychological phenomenon. The production of sentences make public our disagreements; it does not create them’ (Jackson & Pettit, 1998, p. 251).

In order to evaluate this claim, we need to look at what happens when these things come apart. In the case of (23) and (24), we assumed that both Mary and John were sincere. What happens if we suspend that assumption? Let us suppose that John was insincere when he asserted that *War and Peace* was written by Dostoyevsky, and that he believed that *War and Peace* was written by Tolstoy all along. In that case, I am inclined to say that Mary and John do not disagree.² The same point holds in the case (24). If either Mary or John turned out to be insincere, they would not agree.

These judgements are admittedly somewhat delicate, and one should be careful about generalising based on such a limited number of examples. Still, I take this to be some evidence that what matters are not what Mary and John are say-

²It is important to separate the question of whether Mary and John disagree from the question of whether Mary disagrees with what John said. If John is being insincere, these things may come apart.

ing, but what they believe. I am therefore going to side with Jackson and Pettit (1998) and assume that agreement and disagreement are phenomena at the level of mental states, not at the level of language. This does not mean that it is a mistake to focus on examples at the level of language. This is useful for eliciting clear judgements about agreement and disagreement as long as we can assume that the parties are sincere. In general, judgements appear to be more robust when there is no conflict between what is going on at the level of language and what is going at the level of mental states.

While I am going to focus on the attitude question, there is a related question concerning the objects of agreement and disagreement. I assume that if we agree or disagree we might sensibly ask what the object, or objects, of our agreement or disagreement is. In the case of (23), an obvious candidate is the proposition that Tolstoy wrote *War and Peace*. This is the proposition that Mary believes to be true and that John believes to be false. Similarly, a good candidate in the case of (24) is the proposition that Kasparov became world chess champion in 1985 since both Mary and John believe that proposition to be true. This raises the question of whether there is anything that can be said in general about the objects of agreement and disagreement. For instance, if it is true that the object of the disagreement in (23) is a proposition, we might wonder whether this is true in general. Let us call this the ‘object question’.

The Object Question: What are the objects of agreement and disagreement?

A lot of recent discussion about agreement and disagreement can be seen as concerning the objects of agreement and disagreement. It has been argued that we should take the objects of agreement and disagreement to be propositions that are true or false only relative to, say, different individuals or different perspectives. Kölbel (2002, 2004), among others, has argued that this allows us to make sense of alleged cases of faultless disagreement, cases of disagreement in which neither party is making a mistake.

However, here I am going to focus on the attitude question and I am only going to address the object question insofar as it is a connection with the attitude question. It is plausible that there is some kind of connection. Indeed, a straightforward way of thinking about this is to take the objects of agreement and disagreement to correspond to the objects of the attitudes that are involved when

we agree or disagree. In that case, if one thinks that agreement and disagreement can involve attitudes that are not attitudes towards propositions, there will also be cases in which the objects of the agreement or disagreement are not a propositions. However, these questions are at most going to play a secondary role in the following discussion.³

2.2 The Doxastic View

In light of the examples that we considered in the previous section, one might wonder whether there is a simple answer to the attitude question. That is to say that agreement and disagreement always involve belief. This might appear to be a natural and plausible suggestion given what was said about Mary and John's disagreement in (23) and Mary and John's agreement in (24).

(23) Harry: Who wrote *War and Peace*?

- a. Mary: *War and Peace* was written by Tolstoy.
- b. John: No, it was written by Dostoyevsky, not Tolstoy.

(24) Harry: When did Kasparov become world chess champion?

- a. Mary: Kasparov became world chess champion in 1985.

³There are two points of clarifications regarding the scope of the present discussion that I want to make at this point. More specifically, I want to set aside certain uses of 'agree' and 'disagree'. The first point of clarification is that I am mainly interested in agreement and disagreement between two or more individuals. Sometimes we say that we agree or disagree with, say, a theory or an opinion. In that case there may not be anyone that we agree or disagree with because there might not be anyone who believes the proposition in question. For the most part, I am going to ignore such uses of 'agree' and 'disagree' and focus on agreement and disagreement between individuals.

The second point of clarification concerns a distinction that Cappelen and Hawthorne (2009) draws between two different uses of 'agree' and 'disagree'. They argue that there is a performative use that denotes an activity and a stative use that denotes a state of a plurality of individuals. It is only the stative use that I am interested in here. Cappelen and Hawthorne observe that if we are dealing with the performative use of 'agree', there is a sense in which I have agreed with you simply if I say 'I agree' in response to something that you say. I therefore worry that the terms of the debate would change considerably if we were talking about the performative use of 'agree' and 'disagree', and I will only be concerned with the states which are denoted by stative uses of 'agree' and 'disagree'.

b. John: Yes, he became world chess champion in 1985.

The idea is to take the initial observations about (23) and (24) at more or less face value and attempt to understand all cases of agreement and disagreement in the same way. This way of answering the attitude question forms the basis of what I will call the ‘doxastic view’. The central thesis of the doxastic view can be stated as follows:

Thesis 1: Whether two parties agree or disagree is always a matter of which propositions they believe.

Insofar as the objects of belief are propositions, and insofar as the objects of agreement and disagreement correspond to the objects of the attitudes involved, the doxastic view also comes with a commitment to the following thesis:

Thesis 2: The objects of agreement and disagreement are propositions.

More importantly, we can also flesh out the view in order to say something about which attitudes the parties need to have towards a proposition in order to agree, and which attitudes the parties need to have towards a proposition in order to disagree. The most straightforward way of doing this is to add the following necessary conditions on agreement and disagreement:

Thesis 3a: Two parties agree only if there is a proposition that both parties believe.

Thesis 3b: Two parties disagree only if there is a proposition such that one party believes that proposition and the other party believes the negation of that proposition.

There are a lot of complications here, but the precise details are not what is important since I am mainly interested in is the basic idea behind the doxastic view, that is, the commitment to viewing agreement and disagreement as a matter of which propositions that we believe.⁴ For instance, if we want to have a primitive attitude of rejection or denial, such that rejecting a proposition cannot be

⁴According to MacFarlane (2007), we have to weaken these necessary conditions if we are dealing with tensed or temporally neutral propositions, propositions that are true or false relative to times. Suppose that Mary at noon believes the tensed proposition that Socrates is sitting and that

understood simply in terms of believing its negation, we can easily make room for that.⁵ However, for the sake of simplicity, I will assume that the doxastic view comes with a commitment to something like these necessary conditions on agreement and disagreement, even though there is considerable room for discussion when it comes to the details.

Looking at recent debates about relativist treatments of expressions such as predicates of taste and epistemic modals, one easily gets the impression that the doxastic view is widely accepted. One can find both proponents and opponents of relativism expressing support for the doxastic view.⁶ For instance, Kölbel, who defends a version of relativism, makes the following claim:

A case of disagreement typically provides both communicators with a reason to conduct a discussion. For in these cases there is a content p such that one communicator believes p and the other believes not- p .
(Kölbel, 2002, p. 96)

On a relatively straightforward reading this sounds like an expression of support for the doxastic view. If two communicators disagree, there is a content or proposition such that one of them believes that proposition and the other believes the negation of that proposition.

MacFarlane (2007) also endorses the doxastic view, or at least something very close to it, while attempting to give a relativist account of disagreement.⁷ He talks

Peter at midnight believes the negation of the tensed proposition that Socrates was sitting 12 hours ago. It certainly seems right to say that they disagree, but there is no proposition such that one of them believes that proposition and the other believes its negation. Instead, MacFarlane proposes that we should talk about suitably related propositions. This means that it is also misleading to talk as if there is a single proposition that is the object of our agreement or disagreement. Similar considerations come up in the case in which the object of the belief is a centred propositions. However, I am going to set these complications aside for the purpose of the following discussion. See also Chapter 1 for relevant discussion.

⁵See e.g. MacFarlane (2007) who talks about disagreement in terms of acceptance or rejection. See also §2.7 for further discussion.

⁶A notable exception is Weatherson (2009) who explicitly rejects the doxastic view. See also Huvenes (forthcoming).

⁷This is not meant to be an accurate description of MacFarlane's current views on disagreement. In personal communication and recent unpublished work MacFarlane indicates that he favours a different view about disagreement that is more similar to the view associated with Stevenson (1937, 1944, 1963). See MacFarlane (ms.).

about disagreement in terms of acceptance and rejection, but he is reluctant to say that it is a necessary condition on disagreement that there is a proposition that one party accepts and the other party rejects. However, his reasons for being cautious mainly have to do with problems involving tensed propositions and the like, and he wants to allow for cases of disagreement in which one party accepts a proposition and another party rejects a suitably related proposition.

To get a fully general necessary condition [for disagreement], we need to say what it is for two propositions to be ‘suitably related’ such that acceptance of one and rejection of the other constitute disagreement. (MacFarlane, 2007, p. 24)

Even if this is strictly speaking a weakening of the thesis that there must a proposition that one party accepts and the other rejects in order for them to disagree, it is clearly still in keeping with the spirit of the doxastic view.

We also find endorsements of the doxastic view among those who oppose relativism. Cappelen and Hawthorne (2009) fall into this category.⁸ In passages like the following it is clear that they mean to endorse something like the doxastic view:

Thus, the natural picture of agreement we have endorsed above, according to which agreement between a pair of individuals fundamentally consists in their acceptance of the same proposition, is called into question. (Cappelen & Hawthorne, 2009, p. 63)

While I want to argue that the doxastic view is untenable as a view of agreement and disagreement in general, I do not want to deny that the view has considerable appeal. If we look at the examples of agreement and disagreement that we started out with, it is easy to see why someone might be tempted by the doxastic view. For instance, in the case of Mary and John’s disagreement in (23), it sounds right to say that the object of their disagreement is the proposition that Tolstoy wrote *War and Peace*. Whereas Mary believes the proposition that Tolstoy wrote *War*

⁸It is also worth noting that Cappelen and Hawthorne (2009) take the thesis that propositions are the objects of agreement and disagreement to be one of the five theses that make up the view they call ‘Simplicity’. They take Simplicity to be a ‘mainstream’ view and they go on to defend it from various objections, including those of relativists like Kölbel and MacFarlane.

and Peace, John believes its negation. As an explanation of why they disagree this might be a bit simple, but something along these lines is surely not far off the mark. Similarly, in the case of Mary and John's agreement in (24), it sounds right to say that the object of their agreement is the proposition that Kasparov became world chess champion in 1985, a proposition that they both believe to be true. Even if the details are not quite right, something like this better be on the right track. Indeed, it is hard to deny that agreement and disagreement is at least sometimes a matter of which propositions we believe.

I do not deny that the doxastic view may appear to be an initially attractive view about agreement and disagreement. In fact, I think that we should give up the doxastic view only if we have good reasons to do so. There is at least some pressure to try to understand potentially problematic cases of agreement and disagreement as involving beliefs if we can do so. Having said that, I think that we do have reasons to give up the doxastic view. There are cases of agreement and disagreement that do not easily fit into the same mould as the cases we started out with. These cases cannot be understood in terms of what the parties believe, at least not without incurring further and potentially problematic commitments.

2.3 A Problem for the Doxastic View

According to the doxastic view, all cases of agreement and disagreement can be understood more or less along the same lines as the simple cases that we considered in §2.1. However, I want to argue that things are not quite as simple as that. Even if the examples of agreement and disagreement that we considered in §2.1 are easily accommodated if we accept the doxastic view, there are other examples of agreement and disagreement that are more difficult to handle. In this section I am going to focus on cases of disagreement that present a challenge for the doxastic view.

Imagine that Mary and John are a couple who want to spend some time together, and that they are trying to figure out what to do. While they are talking about how they could spend their evening, the exchange in (26) takes place.

- (25) a. Mary: I want to go to the movies.
b. John: No, I would rather just stay at home.

There certainly seems to be a sense in which Mary and John disagree in the dialogue above. For instance, consider how natural it is for a third party to report that they disagree by using (26).⁹

(26) Mary and John disagree about what to do.

I take it that there is nothing particularly strange or out of the ordinary about Mary and John's disagreement in (25) insofar as it is easy to come up with similar examples. Mary and John could be talking about their upcoming holiday and disagree about where to go, or they might be talking about where they want to go for dinner and disagree about whether to go for Indian or Thai food. Moreover, as we should expect, there are analogous examples of agreement. Consider, for instance the following dialogue between Mary and John:

(27) a. Mary: I want to go to the movies.
b. John: Yes, I would like that too.

Just as Mary and John appear to disagree in (25), they appear to agree in (27). In the following discussion I will be focusing on the disagreement between Mary and John in (25), but it is important to bear in mind that there are analogous cases of agreement that we also need to make sense of.

The question is whether we can properly capture the disagreement between Mary and John in (25) within the framework provided by the doxastic view. If the doxastic view is correct, the disagreement between Mary and John is to be understood in terms of the propositions that they believe. There must be a proposition such that one of them believes that proposition while the other believes its negation. But what is the relevant proposition in this case? In the cases that we considered in §2.1, it was a relatively straightforward matter to identify the relevant propositions. However, in this case, matters appear to be much more complicated.

⁹It should be noted that examples like (25) are by no means new in the philosophical literature. Stevenson (1937, p. 27) offers the following example:

A: Let's go to the cinema tonight.

B: I don't want to do that. Let's go to the symphony.

Stevenson takes this to be an instance of what he calls 'disagreement in interest' and distinguishes it from what he calls 'disagreement in belief'.

What should someone who wants to defend the doxastic view say about an example like (25)? There is a natural response that is available to someone who wants to defend the doxastic view and that is to say that Mary and John disagree about whether they should go to the movies. In other words, the suggestion is that the object of their disagreement is the proposition that Mary and John should go to the movies. Whereas Mary believes that proposition, John believes its negation. This may sound like a plausible suggestion, but I want to argue that there are good reasons to think that it is ultimately problematic. Moreover, what I have to say about this specific proposal can hopefully also be extended to other similar responses on behalf of the doxastic view. As far as I can see, nothing turns on whether we are talking about what Mary and John should do, or what they ought to do, or what they have reasons to do, and so forth.

There are two main worries about this response on behalf of the doxastic view. The first worry concerns the sense of ‘should’ that we are talking about.¹⁰ What exactly is the proposition that is the object of Mary and John’s disagreement? The basic problem is to come up with a sense of ‘should’ such that it is both plausible that Mary believes the proposition that they should go to the movies and that John believes the negation of that proposition. The worry is that proposition is going to be too weak and they can both believe it, or it is going to be too strong and neither of them has to believe it. Let us call this the ‘too strong or too weak’ worry.

It does not help us to capture the disagreement between Mary and John if Mary only believes the proposition that they should go to the movies in view of her preferences. John does not have to believe that proposition to be false, he might very well believe that it is true. What John presumably believes to be false is a different proposition, namely the proposition that they should go to the movies in view of his preferences. But this is again not a proposition that Mary has to believe to be true. Mary might very well believe that they should not go to the movies in view of John’s preferences.

Perhaps it makes more sense to say that Mary and John disagree about whether

¹⁰I am assuming a view of deontic modals like ‘should’ and ‘ought’ according to which they are effectively treated as restricted quantifiers over possible worlds. See e.g. Kratzer (1977). However, the point probably goes through on a different view of deontic modals as long they are context-dependent in the right way.

they should go to the movies in view of their joint preferences or in view of what is good for both them. The suggestion would be that Mary not only believes the proposition that they should go to the movies in view of her preferences, she also believes the proposition that they should go to the movies in view of both her and John's preferences.

I do not want to deny that Mary and John could have a disagreement like this. The question is whether this is the only way we can understand Mary and John as having a genuine disagreement. Does Mary have to believe that they should go to the movies in view of both her and John's preferences in order for there to be a disagreement between her and John? There are good reasons to think that this is not the case. Mary might know that John does not want to go to the movies. In that case it seems a bit strange to say that she still has to believe the proposition that they should go to the movies in view of both her and John's preferences. While she might think that her preferences somehow outweigh John's preferences, she certainly does not have to think that.¹¹ However, even if Mary does not think anything like that, we are left with the impression that they disagree.

This discussion is not meant to be exhaustive in the sense that I aim to talk about every conceivable way of trying to make sense of the disagreement between Mary and John. However, I hope that the general point is reasonably clear. It is not easy to identify a sense of 'should' such that it is reasonable to assume both that Mary believes the proposition that they should go to the movies, and that John believes its negation. Either we are left with a proposition that is too weak, or we are left with a proposition that is too strong.

I also want to raise a second worry with the response that Mary and John disagree about the proposition that they should go to the movies. The worry is based on the observation that it is not clear whether it makes much sense to say that either Mary or John is either right or wrong in this case. Let us call this the 'no right or wrong' worry. Manley (2009) makes a point along these lines while commenting on the following example:

(28) a. Christine: Let's go to the beach today.

¹¹Moreover, as Jonathan Ichikawa (p.c.) has pointed out to me, we can add more parties to the conversation. In that case it becomes even more strange to say that Mary has to think that her preferences outweigh the preferences of all the other parties.

b. Melissa: No, let's go downtown instead.

He argues that while there is a sense in which Christine and Melissa disagree, there is no sense in which either of them is right or wrong. There is something strange about saying that one of them is right and the other is wrong. To the extent that this is right in the case of Christine and Melissa's disagreement in (28), it is presumably also right in the case of Mary and John's disagreement in (25).

This raises a worry concerning the suggestion that Mary and John disagree about the proposition that they should go to the movies. If Mary believes the proposition that they should go to the movies and John believes its negation, we would expect one of them to be right and the other to be wrong. But this is not the impression that we are left with. Indeed, it does not look like either Mary or John is either right or wrong in this case.

While I have only considered one response on behalf of the doxastic view, namely the response that Mary and John disagree because Mary believes the proposition that they should go to the movies and John believes its negation, the worries that I have raised are hopefully sufficiently general to cast doubts on the general project of trying to provide a response along these lines. Moreover, it is not clear what an alternative response on behalf of the doxastic view would look like if it was not something along these lines.

If what I have said here is on the right track, cases like (25) constitute a challenge for the doxastic view. That does not mean that we cannot find similar cases in which the parties disagree about, say, what they should do. Those will be cases in which there is a proposition that one party believes to be true and the other party believes to be false. What I am arguing against is the view that all cases of disagreement are like that, and cases like (25) seem to present a problem for this view.

Furthermore, cases like (25) cease to be a problem once we accept that disagreement may involve conflicting non-doxastic attitudes like desires or preferences. A straightforward way of making sense of Mary and John's disagreement is to say that there is a conflict in what they want. What is going on in the case of (25) is simply that Mary wants it to be the case that they are going to the movies, and John wants it to be the case that they are not going to the movies. From this

point of view, there is no real problem here.

2.4 Further Problems for the Doxastic View

In the previous section it was suggested that some cases of disagreement may be characterised as involving conflicting desires rather than conflicting beliefs. In this section, I will argue that two parties can also disagree as a result of having different preferences, say, when it comes to food or humour. For instance, we might want to say that there is a sense in which John and Mary disagree on the basis of (29a) and (29b).

- (29) a. John: I like Dave's curry.
b. Mary: I dislike Dave's curry.

It sounds like (30) is a natural way of reporting their disagreement.

- (30) John and Mary disagree about Dave's curry.

Some philosophers might feel a bit squeamish about saying that John and Mary disagree, but I am not alone in drawing attention to this kind of disagreement.¹² For instance, Weatherson (2009) makes the following observation:

[...] two people can disagree without there being any proposition that one says is true and the other is false. (This should be familiar from debates about non-cognitivism in ethics.) If A says 'I like ice cream' and B says 'I don't like ice cream', then there is a natural sense in which they are disagreeing. (Weatherson, 2009, p. 347)

Schroeder (2008), echoing Stevenson (1937), also seems to acknowledge that there could be a disagreement like this while discussing speaker subjectivism as a view in metaethics.

Speaker subjectivists can take advantage of a move made by Stevenson (1937) and insist that 'murder is wrong' and 'murder is not wrong' are not like 'I'm in Seattle' and 'I'm in New York'; rather,

¹²It is worth point out that not everyone agrees that cases like (29) are cases of disagreement. See e.g. Jackson (2008, p. 81).

they can say, they are more like ‘I’m for the Mariners’ and ‘I’m for the Yankees’. Two people who say the latter, at least sincerely, *are* in a *kind* of disagreement—disagreement in attitude. (Schroeder, 2008, p. 17)

It is not hard to see why examples like (29) present a *prima facie* problem for the doxastic view. The challenge facing the doxastic view is once again to identify a proposition that John believes to be true and Mary believes to be false, but it is far from clear that there are any suitable candidates. It is not enough that John believes the proposition that he likes Dave’s curry and that Mary believes the proposition that she dislikes Dave’s curry. John might believe the proposition that Mary dislikes Dave’s curry and Mary might believe the proposition that John likes Dave’s curry. This does not explain why they disagree. If the doxastic view is correct, there must be some other proposition such that one of them believes that proposition and the other believes its negation.

The natural response is to argue that the object of the disagreement is the proposition that Dave’s curry tastes good. This response is similar to the response on behalf of the doxastic view that was discussed in the previous section, and it turns out to be problematic for similar reasons. In the previous section I raised two worries, the ‘too strong or too weak’ worry and the ‘no right or wrong’ worry, concerning this way of attempting to explain Mary and John’s disagreement about what to do, and I will argue that there are similar problems with the response that is under discussion in this section.

The first worry is analogous to the ‘too strong or too weak’ worry. It is hard to identify a proposition such that it is plausible that John believes that proposition and Mary believes its negation. It is fair to assume that John believes the proposition that Dave’s curry tastes good to him since he is willing to assert that he likes the curry, but Mary does not have to believe that proposition to be false. What she believes, is the proposition that Dave’s curry does not taste good to her. But this is again something that John can believe too.

On the other hand, if we take the relevant proposition to be the proposition that Dave’s curry tastes good to both of them, it is no longer clear that John has to believe that proposition. He does not even have to have a view about whether

Dave's curry tastes good to Mary or not.¹³ Taking the object of their disagreement to be something like the proposition that Dave's curry tastes good to members of the wider community, does not help. Even if we assume that there is a community to which both John and Mary belong, there is not much of a basis for thinking that John has to believe the proposition that Dave's curry tastes good to members of this community. After all, he is only asserting that he likes the curry.¹⁴

The point is not that John and Mary could not have a disagreement like this. The point is that there does not seem to be strong case for thinking that the disagreement in (29) has to be understood in this way. Moreover, as far as I can see, the judgement that John and Mary disagree is not contingent on John believing anything stronger than the proposition that Dave's curry tastes good to him, and that is not enough to make sense of the disagreement if the doxastic view is correct.

There is also a worry analogous to the 'no right or wrong' worry. If John and Mary's disagreement can be explained by there being a proposition such that John believes that proposition and Mary believes its negation, we should expect one of them to be right and one of them to be wrong. Typically, when someone believes a proposition and someone else believes the negation of that proposition, we expect one of them to be getting it wrong. But it does not look like John or Mary have to be making a mistake in (29). There is a sense in which neither John nor Mary is getting it either right or wrong in this case. This suggests that we should look for a different way of handling cases like this, one which does not require there to be a proposition such that one of them believes that proposition and the other believes its negation.

As in the previous section, these considerations do not amount to anything like a decisive refutation. For one thing, I do not pretend that I have explored

¹³In fact, we do not even have to assume that John and Mary are participants in the same conversation, or even that they know each other, to generate a sense of disagreement. If they do not even know each other, it would be extremely odd for John to have to have a view about whether the curry tastes good to Mary.

¹⁴It is important to note that what is at stake here is not whether matters of taste are somehow more or less objective. Even if it turns out that matters of taste are perfectly objective, whatever that may involve, that does not make it easier to explain John and Mary's disagreement in (29). If anything, it would be even harder to motivate the claim that John believes the proposition that Dave's curry tastes good if that was understood in a perfectly objective sense.

every possible way of making sense of the idea that there is a proposition that John believes to be true and Mary believes to be false and that this is why they disagree. However, I hope that I have done enough to show that there are general worries about this line of response that go beyond my criticism of the specific proposals. Instead of trying to look for a proposition such that one of them believes that proposition and the other believes its negation, we can deal with cases like (29) by taking them to involve a conflict of non-doxastic attitudes. We simply take appearances at face value and say that John and Mary disagree because he likes Dave's curry and she does not like it. What we have is a case of conflicting non-doxastic attitudes and not a case of conflicting beliefs.

2.5 Relativism to the Rescue?

While the considerations that I have put forward in the previous sections put significant pressure on the doxastic view, someone might object that I have been too quick. There is a possible response on behalf of the doxastic view that I have been ignoring so far. Someone who wants to hold on to the doxastic view could adopt a form of relativism in an attempt to get around the objections that I have raised. Relativism has recently attracted a lot of attention as a means of explaining allegedly puzzling disagreement data involving, among other things, predicates of taste and epistemic modals.¹⁵

Different relativists have articulated their views in different ways. For the sake of simplicity, I am going to more or less follow Egan (2007) in taking relativism to be the view that there are propositions that have different truth-values relative to different individuals, but what I am about to say can be adapted to a different versions of relativism.¹⁶ Let us consider the case of Mary and John's disagreement in (25) in light of this view.¹⁷ The idea is that the object of their disagreement is the proposition that Mary and John should go to the movies. However, that

¹⁵See e.g. Kölbel (2002, 2004), Egan et al. (2005), Lasersohn (2005), MacFarlane (2005b, 2007), Egan (2007), and Stephenson (2007).

¹⁶Egan (2007) formulates his relativist account in terms of there being propositions that determine a truth-value relative to a centred world, a <world, time, individual> triple, and not just a possible world. See Chapter 1 for further discussion of relativism.

¹⁷For a more worked out relativist treatment of predicates of taste such as 'tasty' that is relevant with respect to the discussion in §2.4, see e.g. Lasersohn (2005) and Stephenson (2007).

proposition is true relative to Mary and false relative to John. For instance, it might be true relative to Mary in virtue of her preferences and false relative to John in virtue of his preferences.¹⁸

This line of response on behalf the doxastic view should not come as a surprise to someone who has observed recent debates about relativism in philosophy of language and linguistics. In fact, the style of argument that I have employed against the doxastic view, is similar to the style of argument that is often used to motivate relativism. It would therefore not be surprising if adopting a relativist position can give us some traction with the ‘too strong or too weak’ worry and the ‘no right or wrong’ worry. Since the proposition that Mary and John should go to the movies is true relative to Mary, it makes sense for her to believe it and since the negation is true relative John it makes sense for him to believe the negation. Moreover, since the proposition that Mary believes is true relative to her and since the proposition that John believes is true relative to him, there is a sense in which neither of them has made a mistake.

All of this is of course contingent on there being a plausible version of relativism available. This is both a specific question regarding the prospects of a relativist treatment of the relevant expressions, and a question regarding the general prospects of relativism. One must both provide a plausible relativist treatment of ‘should’ and show that such a relativist account can secure genuine disagreement. These tasks are not trivial.

I do not have much to offer in the way of objections to relativism at this point, either generally or with regard to the prospects of using relativism to provide a defence of the doxastic view. A lot has been said on this subject by others, and I am not going to repeat it here.¹⁹ I regard it as an open question whether it is ultimately possible to offer a plausible account of disagreement along the lines that relativists propose. What I want to emphasise is that this way of defending the doxastic view comes with significant commitments. If one is forced to adopt a relativist position in order to hold on to the doxastic view, that might turn out to be a high price to pay. There is at least something to be said for exploring

¹⁸For a more sophisticated discussion of deontic modals, see e.g. Kolodny and MacFarlane (2010).

¹⁹See e.g. Stojanovic (2007), Moruzzi (2008), Rosenkranz (2008), and Francén (2010) for some recent criticisms of relativist accounts of disagreement.

other options.

What it is also important to recognise, is that the relationship between relativism and the doxastic view is not just a one way relationship. Relativism offers a way of responding to the arguments against the doxastic view. But relativism is also supposed to be motivated by considerations involving disagreement. We have seen that Kölbel (2002, 2004) and MacFarlane (2007) endorse something like the doxastic view. Without the doxastic view it becomes more difficult to motivate relativism by appealing to considerations involving disagreement. If disagreement can involve a wide range of attitudes, there is worry that it becomes more difficult to ensure that the relevant cases should be understood in the way that relativists want them to be understood.²⁰

2.6 Disagreement in Attitude

The picture of agreement and disagreement that I want to defend is closely tied to a way of thinking about agreement and disagreement that goes back to Stevenson (1937, 1944, 1963). Stevenson distinguished between what he called ‘disagreement in belief’ and what he called ‘disagreement in attitude’ or ‘disagreement in interest’. There is disagreement in belief when there is a proposition such that one party believes that proposition and another party believes its negation. This is the kind of disagreement we find in (23). Disagreement in attitude contrasts with disagreement in belief insofar as it involves a conflict of attitudes as opposed to a conflict of beliefs.

Two men will be said to disagree in attitude when they have opposed attitudes to the same object—one approving of it, for instance, and the other disapproving of it—and when at least one of them has a motive for altering or calling into question the attitude of the other.
(Stevenson, 1944, p. 3)²¹

It is this way of thinking about disagreement, and the notion of ‘disagreement in attitude’ in particular, that I am interested in. According to this picture, agree-

²⁰See Chapter 3 for further discussion.

²¹See also Stevenson (1937, 1963).

ment and disagreement do not always have to involve belief. Sometimes agreement and disagreement involve other attitudes, such as approval and disapproval.

This way of thinking about agreement and disagreement can also be found in the works of contemporary expressivists like Blackburn (1984, 1998) and Gibbard (2003).²² For instance, consider the following passage from Blackburn (1998):

If I am minded to permit smoking in our house, and my wife is minded to forbid it, we do disagree. Only one of these practical attitudes can be implemented, and I am for one, and she is for the other. When we discuss ethics with each other, we are typically talking about ‘our house’, or in other words practical issues on which we want to coordinate, or have to coordinate. In that case difference of attitude means disagreement, just as surely as disagreement in belief does. [...] the typical, default, position is that difference in attitude is treated as disagreement. (Blackburn, 1998, p. 69)

In this passage, Blackburn appears to be endorsing a view of disagreement that is similar to the view proposed by Stevenson. But even if this way of thinking about agreement and disagreement might be attractive from an expressivist point of view, that does not mean one is committed to some kind of expressivist view if one thinks about agreement and disagreement in this way. The picture of agreement and disagreement that I want to defend, is compatible with the denial of any sort of expressivism as a thesis about, say, the workings of normative thought and talk. Thinking that agreement and disagreement can involve attitudes other than belief, does not entail any sort of expressivist thesis. What matters for our purposes are the attitudes of the parties, not how these attitudes are expressed in language.

Even though I take Stevenson’s notion of ‘disagreement in attitude’ as my starting point, I want to point out that my preferred way of thinking about these

²²Gibbard (2003) focuses on the notion of ‘disagreement in plan’ and contrasts this with Stevenson’s notion of ‘disagreement in attitude’. However, insofar as disagreement in plan also contrasts with disagreement in belief, I do not think that there is any need to think that there is a fundamental conflict here. As far as I can see, there is room for both disagreement in plan and disagreement in attitude as Gibbard construes these notions.

matters differs from Stevenson's in certain respects. The main point of departure from the picture that Stevenson presents, is that I do not take it to be an essential part of the view that there are two different senses of 'agree' and 'disagree'. If we understand 'attitude' in a broader sense that also includes belief, we can think of all agreement as involving converging attitudes and all disagreement as involving conflicting attitudes. Having conflicting beliefs is just one way of having conflicting attitudes. In effect, the proposal is that all agreement and disagreement is agreement and disagreement in attitude, but with 'attitude' being understood in a sufficiently broad sense that also includes belief.

It is not entirely clear in what way Stevenson thought that we needed to distinguish between disagreement in belief and disagreement in attitude, and it is therefore not clear how substantial these differences are. In any case, the point is not that 'agree' and 'disagree' cannot exhibit some form of context-dependence or ambiguity. Whether 'agree' and 'disagree' are context-dependent or ambiguous is an empirical question about the semantics of English, and one would have to examine the empirical evidence in order to settle that question. The point is that we do not have to think there are two different senses of 'agree' and 'disagree' in order to allow for cases of agreement and disagreement that involve attitudes other than belief.²³

With these clarifications in mind, I want to put forward the following two theses:

Thesis 1*: Two parties agree just in case they have converging attitudes.

Thesis 2*: Two parties disagree just in case they have conflicting attitudes.

While this tells us something about agreement and disagreement, agreement and disagreement involve converging and conflicting attitudes respectively, it still leaves a lot of questions unanswered. What is it for attitudes to conflict or converge? What are the attitudes in question? Are we dealing with a wide range of attitudes or only a few special attitudes? At the very least these theses have to be

²³Another potential point of departure from Stevenson's account is that he took it to be a necessary condition on disagreement that one of the parties has to be motivated to change or challenge the attitude of the other party. It is not clear how well-motivated this condition is, and I take it to be an open question whether it is needed. However, note that Stevenson (1963) also took this to be a necessary condition for disagreement in belief and not just for disagreement in attitude.

supplemented with the thesis that there are cases of conflicting and converging attitudes that are not cases of conflicting and converging beliefs.

Pluralism: There is convergence and conflict among a wide range of attitudes that includes both doxastic and non-doxastic attitudes.

This thesis, together with Thesis 1* and Thesis 2*, is the central thesis that I want to defend. It captures the idea that agreement and disagreement may involve other attitudes than belief. What are the relevant attitudes apart from beliefs? Stevenson (1944, p. 3) mentions purposes, aspirations, wants, preferences, and desires. Other potential examples are liking, admiring, hoping, and so forth. For instance, if someone desires or wants something to be the case and someone else wants it not to be the case, there is a sense in which their attitudes conflict. Similarly, two parties can be said to have conflicting attitudes if there is something that one of them likes and the other dislikes. In these cases it sounds quite natural to say that there is some sort of disagreement between the parties.

In light of this, it should be clear that the view should not have much trouble with the cases of disagreement that we have discussed so far. There is nothing that prevents this view from dealing with the cases in §2.1 in more or less the same way as the doxastic view. Moreover, if we can think about disagreement in terms of what the parties want or like, the cases in §2.3 and §2.4 should also be easier to deal with. For instance, it seems plausible to say that the disagreement between Mary and John in (25) is a matter of Mary and John having conflicting desires or preferences.

- (25) a. Mary: I want to go to the movies.
b. John: No, I would rather just stay at home.

Examples like this are useful as a starting point for getting a better grip on what conflicting and converging attitudes are, but it would be nice to be able to say something more general, something that goes beyond particular examples. Is there anything general to be said about what it is for attitudes to conflict or converge? Again, we should not expect too much. After all, we are not aiming for a fully reductive account of agreement and disagreement.

There does not have to be anything particularly mysterious about talking about people as having conflicting or converging attitudes. For one thing, we

routinely compare the attitudes of different people. We are often interested in whether other people have the same attitudes as we do. On many occasions, it also makes sense to ask whether our attitudes are in conflict or whether they are compatible.

In light of this, I think that it is possible to say something about conflicting and converging attitudes that might help us to get a better grip on the view of agreement and disagreement under discussion. In the next section I will try to say something more about conflicting attitudes with the goal of giving us a better grip on this notion. In doing so, I will be focusing on conflicting attitudes instead of converging attitudes. Once we have a better grip on conflicting attitudes this will hopefully also give us a better grip on converging attitudes as well.

2.7 Conflicting Attitudes

There are broadly speaking two ways in which two parties can be said to have conflicting attitudes.²⁴ The first is that there is a proposition such that one of the parties has a certain attitude towards that proposition, and the other party has the same attitude towards the negation of that proposition.²⁵ For instance, when two parties have conflicting beliefs there is a typically a proposition such that one party believes that proposition, and the other party believes its negation. Along similar lines, there might be a proposition such that one wants it to be true, and the other party wants its negation to be true.

The second way in which two parties can have conflicting attitudes, is that they have different and conflicting attitudes towards the same object. A possible example of this would be if there is something that one party likes and the other party dislikes. It is worth pointing out that, depending on one's views about the attitudes in question, the relevant object does not have to be a proposition. Insofar as this is the case, there is no reason to think that the objects of agreement

²⁴This distinction is similar to the distinction that Schroeder (2008) draws between what he calls 'A-type' and 'B-type' inconsistency. Schroeder argues that there are no uncontroversial cases of 'B-type' inconsistency. That may be true if we are talking about inconsistency, but it does mean that it has to be true if we are talking about disagreement.

²⁵For the purpose of the following discussion I will assume that propositions are true or false *simpliciter* and not relative to any parameters. For instance, if there are propositions that have different truth-values relative to different times, this would introduce further complications.

and disagreement will always be propositions, contrary to what one might expect if one is only paying attention to cases like (23) and (24).²⁶

In certain respects, the first way in two parties can have conflicting attitudes is the most familiar. However, there are reasons for thinking that all cases of conflicting attitudes should ultimately be understood in terms of the parties having different and conflicting attitudes towards the same object. That is true even in the case of belief. Consider the beliefs of someone who is a dialethist and believes that the liar sentence is true and that the liar sentence is not true. Suppose that she encounters another dialethist with exactly the same beliefs. If it were sufficient for having conflicting attitudes that one party believes a proposition and another party believes its negation, we would expect them to disagree. But that does not seem right. One way of avoiding this problem is to introduce a primitive attitude of rejection or disbelief. Two parties count as having conflicting attitudes if there is a proposition that one party believes and the other rejects. It is typically safe to assume that someone who believes the negation of a proposition also rejects it, but occasionally it is necessary to keep these notions apart.²⁷

For the purpose of the following discussion, I will assume that we can ignore these complications and treat two parties as having conflicting attitudes if there is a proposition such that one party believes that proposition, and the other party believes its negation. More generally, I am going to continue to assume that there are two different ways of having conflicting attitudes. There can be an object such that one party has an attitude towards that object, and the other party has a different and conflicting attitude towards the same object. But there can also be a proposition such that one party has an attitude towards that proposition,

²⁶It is not sufficient for two parties to have conflicting attitudes if one party has an attitude and the other party lacks that attitude. For instance, this may happen if one of the parties may lack the relevant concepts. Suppose that Mary likes her father's cooking. If John has never tasted Mary's father's cooking, he will presumably not share that attitude. However, that does not mean that Mary and John have conflicting attitudes in the relevant sense. In order for Mary and John to have conflicting attitudes, John would have to dislike Mary's father's cooking.

²⁷One might object that it is a problem with dialethism that it severs the connection between negation and disagreement. Since I have no inclination to accept dialethism, I am not going to take a stand on that. However, one does not have to be dialethist to think that two dialethists should not count as having a disagreement in virtue of their beliefs about the liar sentence.

and the other party has the same attitude towards its negation. By doing this I am actually making matters more difficult for myself by introducing further complications and making the case of conflicting beliefs seem more special than it is.

In any case, it is important to see that a difference in attitudes is not always sufficient to give us a conflict of attitudes. If someone supposes that there is an odd number of atoms in the universe, and someone else supposes that there is not an odd number of atoms in the universe, we do not want to say that their attitudes conflict or that they disagree. However, there is a difference in attitude. The same is true of wondering. It is arguably also true of imagining. We do not want to say that two people disagree just because one of them is imagining that there is an odd number of atoms in the universe while the other is imagining that there is not an odd number of atoms in the universe. If this is right, having conflicting attitudes cannot just be a matter of having different attitudes. This assumes that a conflict of attitudes can be a matter of there being a proposition such that one party has a certain attitude towards that proposition, and the other party has the same attitude towards the negation of that proposition. But insofar as that assumption has already been granted, we are faced with the question of how to distinguish between cases involving wondering and imagining, and cases involving wanting and believing.

The question is when a difference in attitudes amounts to a conflict of attitudes. One way of getting some traction with this question is to use intrapersonal cases of conflicting attitudes as a starting point, and attempt to use them to get a better grip on the interpersonal cases.²⁸ That is going to be strategy that I will pursue here. This is relatively easy to see in cases in which there is a conflict among one's beliefs. If one believes a proposition, it is clear that there would be conflict if one also were to believe its negation. However, I take it that there can also be conflicts among one's non-doxastic attitudes. There could for instance be a conflict among one's preferences or desires. If one likes and dislikes chilli at the same time, this will typically constitute a conflict of attitudes.²⁹ Similarly,

²⁸MacFarlane (ms.) also appears to be suggesting something similar.

²⁹Saying that both liking and disliking something at the same time constitutes a conflict of attitudes is perhaps a bit simplistic. It is possible to like something in one respect and dislike it in another respect without any real conflict. For instance, one can like the smell of the coffee and at

there is typically a conflict among one's attitudes if one both wants something to be the case, and at the same time wants it not to be the case.

We do not have to look at actual conflicts. Even if it turns out that actual conflicts of this sort are rare, we can recognise potential conflicts. For instance, if one wants it to be the case that the number of atoms in the universe is odd, one can recognise that there would be a conflict if one also were to want it to be the case that the number of atoms in the universe is not odd. We can also look at the difference between our current attitudes and our previous attitudes. Suppose that Mary used to like chilli, but that she no longer likes it. In that case, her current attitudes are in conflict with her previous attitudes. This suggests that (31) can be used as a guide to whether we have a conflict of attitudes at the intrapersonal level.

- (31) If there is a conflict between two attitudes at the intrapersonal level, those attitudes are also in conflict at the interpersonal level.

In the case of wondering and supposing we typically do not find any such conflicts among our own attitudes.³⁰ This is perhaps easiest to see in the case of wondering. When someone wonders whether a proposition is true, it is entirely natural for them to also wonder whether the negation is true. However, the point also applies in the case in which one is supposing that something is the case. For instance, I can suppose that both a proposition and its negation is true in order to see what follows. I might just be interested in seeing what follows from a contradiction. This does not mean that there is any kind of conflict among my attitudes. Insofar as we do not find conflicts involving these attitudes on the intrapersonal level, it is not surprising that we do not find them at the interpersonal level either.

For the purpose of this discussion, 'conflict' should be understood in a relatively broad sense. For instance, it does not have to entail that it is psychologically impossible to have both attitudes at the same time. Even in the case of conflicting beliefs, it is not clear that this is the case.

the same time dislike its taste.

³⁰The case of imagining is more complicated. However, there is a sense in which there is no conflict between imagining that something is the case, and imagining that something is not the case.

It is also doubtful whether the relevant notion of ‘conflicting attitudes’ should be cashed out in terms of which attitudes it is rationality permissible to have.³¹ Again, the worry is that this would be too restrictive.³² Perhaps there is a sense in which it is not rational to both like and dislike something, but that is not obvious. At the very least this would introduce some highly contentious issues about rationality that I would prefer to avoid for the purpose of this discussion.

A more promising idea is that there is a conflict between two attitudes just in case having both attitudes at the same time tends to lead to problems. For instance, if one has a conflict among one’s desires or preferences one might have trouble figuring out what to do. One might also find oneself in a situation in which it is impossible for one’s desires or preferences to be satisfied. In the case of conflicting beliefs, the problem could be that one’s beliefs could not both be true. One would be guaranteed to end up with a false belief. Moreover, there are no such problems associated with supposing or imagining. It is not the case that imagining a proposition and its negation leads to problems like this. As a guide to whether a difference in attitude counts as a conflict of attitudes, this seems to be promising.

Finally, I want to stress that there are lots of problems involving agreement and disagreement that are problems for everyone regardless of how they answer the attitude question. For instance, how should we think about conflicting and converging attitudes if the objects of the relevant attitudes are centred propositions, propositions that vary in truth-value across individuals? Does it make a difference with respect to whether two parties have conflicting or converging attitudes if they are thinking about the relevant objects under different guises or modes of presentation? These are interesting questions, but they are not questions specifically for someone who rejects the doxastic view. There is no reason to think that someone who endorses the doxastic view will have an easier time answering questions like this.

The question that I am mainly interested in is the attitude question. In par-

³¹Dreier (2009) suggests something along these lines.

³²It is also important to make sure that the relevant condition is not too permissive. It cannot be sufficient for there to be a conflict of attitudes that one cannot rationally have both attitudes. That condition would be satisfied if it was not rational to have one of the attitudes. Instead, the relevant condition would have to be that one would fail to be rational solely in virtue of having those attitudes.

ticular, I want to argue that we should think about agreement and disagreement as potentially involving a wide range of attitudes, and not just belief. I have therefore focused on issues that are fairly closely connected to these issues. In this section I have tried to say something about how we can make sense of what it takes to have conflicting attitudes, when we allow that there can be conflicts among attitudes other than belief. As I have said, I do not think that this notion is particularly mysterious. However, to the extent that someone finds the notion of ‘conflicting attitudes’ problematic, I hope that what I have said can give them a better grip on this notions.

2.8 Faultless Disagreement

A potential benefit and a potential worry with having a more pluralist view of disagreement, is that it leaves room for faultless disagreement. It is not entirely clear what faultless disagreement amounts to because there are many ways in which someone can be said to be wrong or making a mistake. In the case of belief, one might for instance be wrong because one believes something false or because one believes something on the basis of inadequate or misleading evidence. In what follows, I will understand faultless disagreement as disagreement that is compatible with neither party making a mistake in a broad sense of ‘making a mistake’. The fact that two parties disagree does not entail that one of them is making any kind of mistake.

The more pluralist picture of disagreement that was developed in §2.6 and §2.7, appears to be compatible with faultless disagreement in this broad sense. If you can have a disagreement as a result of having, say, conflicting desires or preferences, there is no reason to think that either party is wrong or at fault or making any sort of mistake. If Mary wants it to be the case that she and John are going to the movies, and John wants it to be the case they stay at home, neither of them has to be making a mistake.³³ In fact, when I was arguing that the doxastic

³³That is not to say that it is impossible for one of them to be making a mistake. Perhaps there are circumstances in which John would be making some kind of mistake by wanting Mary and him to stay at home. Nothing that I have said is inconsistent with this being a genuine possibility. However, the stronger claim that this is true in general whenever someone is having a disagreement like the one that Mary and John are having, looks much less plausible.

view has a hard time dealing with cases like this, I argued that we do not want to be forced to say that either Mary or John is making a mistake. Adopting a more pluralist conception of disagreement allows us to avoid saying precisely that. From this point of view, it is a virtue of this sort of account that it is compatible with faultless disagreement.

However, it may also be argued that it is also a vice to predict that there are cases of faultless disagreement. While proponents of a relativist semantics like Kölbel (2002, 2004) and Lasersohn (2005) have taken it to be an advantage of their views that they can secure faultless disagreement, this is far from universally accepted. In fact, the view that there is faultless disagreement has been met with considerable scepticism. For instance, Glanzberg offers the following remarks on the topic of faultless disagreement:

Lasersohn, and a number of other contemporary relativists, point out that their notion of relative truth offers a notion of ‘faultless disagreement’, where two utterances express disagreement, even though neither is incorrect (cf. Kölbel 2002). From a traditional, non-relativist, point of view, this idea is *prima facie* absurd: if two propositions express disagreement, one must fail to be correct. [...] My own inclination is to side with the traditional view, and reject the notion of faultless disagreement as absurd. (Glanzberg, 2007, p. 16)

I take it that the sort of scepticism about faultless disagreement that Glanzberg is expressing is not uncommon.³⁴ But even if one is sceptical about the notion of faultless disagreement that some relativists are trying to secure, that does not necessarily translate into scepticism about the sort of faultless disagreement that one gets if one thinks that disagreement may involve non-doxastic attitudes. It is not clear what kind of disagreement Glanzberg has in mind when he is talking about faultless disagreement. His scepticism might be appropriate insofar as he is talking about cases of disagreement in which the parties are believing contradictory propositions, and perhaps that is how we should interpret him when he talks about two propositions expressing disagreement. The assumption that

³⁴See e.g. Stojanovic (2007), Moruzzi (2008), and Rosenkranz (2008) for various criticisms of relativist attempts to make sense of faultless disagreement.

faultless disagreement must involve contradictory beliefs is made even more explicit by Kölbel (2002, 2004). For instance, Kölbel makes the following claim about faultless disagreement:

There are disagreements without error, or in other words, some propositions are not objective. However, minimal constraints on truth show that if it is true that p , then it is not true that not- p , and if it is true that not- p , then it is not true that p . So if one thinker believes that p and another believes that not- p , one of them makes the mistake of believing a proposition that is not true. The only way to allow faultless disagreement is therefore to relativise truth to perspectives: one disputant's belief is true in his or her own perspective, and the other disputant's contradictory belief is true in his or her own perspective. (Kölbel, 2002, p. 100)

Given the way he thinks about disagreement, it is not strange that Kölbel comes to the conclusion that faultless disagreement requires truth to be relative to perspectives. He wants to be able to make sense of cases in which two parties have contradictory beliefs without either of them being wrong. In the following definition of 'faultless disagreement' from Kölbel (2004), he makes it explicit that the notion of faultless disagreement he is talking about requires the parties to have contradictory beliefs:

A faultless disagreement is a situation where there is a thinker A, a thinker B, and a proposition (content of judgement) p , such that:

- (a) A believes (judges) that p and B believes (judges) that not- p
- (b) Neither A nor B has made a mistake (is at fault).

(Kölbel, 2004, p. 53-54)

Kölbel goes on to claim that 'most people have a healthy pre-theoretic intuition that there can be and are faultless disagreements in this sense'. That may or may not be true. The point is that one might well reject this sort of faultless disagreement while accepting faultless disagreement involving conflicting desires or preferences. Someone who supports a more pluralist conception of agreement and disagreement can grant that there is something *prima facie* odd about

two parties having a disagreement as result of there being a proposition such that one of them believes that proposition and the other believes its negation without either of them being mistaken. But that does not necessarily mean that there is something correspondingly odd about a case of faultless disagreement involving conflicting desires or preferences.

It is useful to look at a more specific objection to faultless disagreement in order to illustrate this point. Rosenkranz (2008) presents a dilemma for relativists like Kölbel who want to make sense of faultless disagreement. Suppose that there is a proposition p such that I assert that p , and that you assert that not- p , and that p is true relative to my perspective, but false relative to your perspective.³⁵ Rosenkranz argues that if I merely present p as true relative to my perspective, and you merely present not- p as true relative to your perspective, we do not really have a disagreement. On the other hand, if we present the propositions in question as true *simpliciter* or relative to every perspective, we do have a disagreement, but it is not faultless because it is a mistake to present p as true *simpliciter* if it is only true relative to my perspective.

I do not want to take a stand on whether the argument succeeds in raising problems for Kölbel's account of faultless disagreement. The initial reaction is that relativists may not want to accept that there is no disagreement if I merely present p as true relative to my perspective, and you present not- p as true relative to your perspective. Alternatively, they might raise questions about the notion of 'presenting a proposition as true'.

In any case, what I want to point out is that this sort of argument does not even present a *prima facie* problem for the kind of faultless disagreement that one gets if one has a more pluralist conception of disagreement. The dilemma, if there is one, arises because relativists like Kölbel relativise propositional truth to perspectives and therefore need to say something about the correctness of assertions and beliefs when the propositions in question have different truth-values relative to different perspectives. But it is hard to see what a corresponding worry about faultless disagreement involving conflicting desires or preferences would look like. The worry concerns the criteria for correctness of assertions

³⁵Rosenkranz (2008) runs his argument mainly in terms of assertion, but I assume that it could also be run in terms of belief. As long as the parties are sincere and not confused in some relevant sense this should not make a difference.

of beliefs, but the faultless disagreement that comes with a pluralist conception of agreement and disagreement involves other kinds of attitudes. If we are only concerned with cases of faultless disagreement that involve attitudes other than belief, the objection would not even get off the ground.

This is of course not to say that faultless disagreement involving conflicting desires or preferences is unproblematic. It would take a lot more arguing to establish that. What needs to be recognised is that the issues surrounding faultless disagreement change if we give up the doxastic view and adopt a more pluralist view of agreement and disagreement. There might be one set of challenges associated with having faultless disagreements in which the parties have conflicting beliefs, and a different set of challenges associated with having faultless disagreements in which the parties have conflicting desires or preferences.

Conclusion

The starting point for this discussion was a question about the attitudes that are involved when we agree or disagree. In an attempt to answer this question, I have outlined what I take to be a promising way of thinking about agreement and disagreement. According to this picture, agreement involves converging attitudes, and disagreement involves conflicting attitudes. Whether we agree or disagree is not just a matter of what we believe. Agreement and disagreement can involve a wide range of attitudes such as wanting, liking, approving, and so forth.

In terms of motivating this picture of agreement and disagreement my goals have been fairly modest. My main concern has been to argue that the doxastic view, which may appear to be an initially attractive view of agreement and disagreement, faces serious problems. While it may be possible to rescue the doxastic view by adopting some kind of relativist position, this is not unproblematic. In particular, it is an open question whether it is possible to make sense of agreement and disagreement in relativist terms. If the relativist account of disagreement is problematic, the cure might well turn out to be worse than the disease. In any case, it is a serious commitment on the part of someone who wants to defend the doxastic view.

Chapter 3

Predicates of Taste

Introduction

Predicates of taste, such as ‘fun’ and ‘tasty’, have received considerable attention in recent debates between contextualists and relativists, with considerations involving disagreement playing a central role.¹ Considerations involving disagreement have been taken to present a problem for contextualist treatments of predicates of taste. My goal is to argue that considerations involving disagreement do not undermine contextualism. To the extent that relativism was supposed to be motivated by contextualists being unable to deal with disagreement, this motivation is lacking. The argument against contextualism rests on a too simple and narrow conception of disagreement that turns out to be problematic once we consider a wider range of cases. If we reject the assumptions about disagreement that the argument rests on, it no longer poses a threat to contextualism.

Supporters of relativism, such as Kölbel (2002, 2004), Lasersohn (2005), MacFarlane (2007), and Stephenson (2007), have argued that contextualist treatments of predicates of taste fail to deliver the right predictions about when speakers disagree. Instead they argue that we should prefer a relativist treatment of predicates of taste. I will argue that this line of argument is not as effective as it initially appears to be. My main concern in this paper is to argue that there are

¹This is a revised version of a paper that is forthcoming in the *Australasian Journal of Philosophy*. See Huvenes (forthcoming).

other cases of disagreement that we need to take into account. The examples I have in mind are examples like (32).

- (32) a. Mary: I like this chilli.
b. John: I disagree, it's too hot for me.

I will argue that examples like this do not call for a relativist treatment. Moreover, once examples like (32) have been taken into account, the cases of disagreement that were supposed to present a problem for contextualism seem to be less surprising from a contextualist point of view.

In §3.1 I outline the difference between contextualist and relativist treatments of predicates of taste. In §3.2 I explain why contextualist treatments allegedly have problems with disagreement. In §3.3 I present some further disagreement data, and argue that the cases that were allegedly problematic from a contextualist point of view are less surprising in light of the new data. In §3.4 I try to assuage some doubts one might have about the data presented in §3.3. In §3.5 I argue that focusing on other disagreement markers does not make the argument against contextualism more effective. In §3.6 I offer some preliminary remarks about the possibility of understanding the disagreement data in terms of conflicting non-doxastic attitudes.

3.1 Contextualism and Relativism

The debate between contextualists and relativists that I am concerned with, is a debate about the semantics of predicates of taste. In terms of characterising these predicates, I will be content with saying that they are formed from adjectives such as 'fun' and 'tasty'. Instead of trying to come up with criteria for identifying the relevant expressions, I will follow Lasersohn (2005) and most of the other participants in this debate and focus on the expressions 'fun' and 'tasty'.

It is common ground between contextualists and relativists that there is some sort of contextual variation in the truth-values of sentences containing 'fun' and 'tasty'. The contextualist and relativist treatments of predicates of taste that I am interested in, take the truth-values of such sentences to depend on whose taste that is relevant in some way or another. However, they differ with respect to how they account for this variation in truth-values. According to the semantic

framework of Kaplan (1989), which I will be assuming for the purpose of this discussion, the content of an expression in a context can be represented as a function from an index, a world-time pair on his view, to an extension.² The extension of a sentence is a truth-value, and the truth-value of a sentence can depend on both the context insofar as it plays a role in determining the content and the index with respect to which we evaluate the content. I will talk about the content of a sentence in a context as a proposition.³

Simply put, contextualism about predicates of taste is the thesis that sentences containing predicates of taste express different propositions in different contexts. According to the kind of contextualist position we are concerned with here, different propositions are expressed in different contexts, depending on whose taste that is relevant. What matters is that the contextual variation in truth-value is explained by different propositions being expressed in different contexts.

A relativist account differs from a contextualist account in that it does not take sentences containing predicates of taste to express different propositions in different contexts. Instead, the relevant sentences express propositions that are true or false relative to some appropriate parameter other than worlds and times. Following Lasersohn (2005), relativists about predicates of taste can take the truth or falsity of sentences containing predicates of taste to be relative to individuals. In terms of modifying the basic semantic framework from Kaplan this means that indices are treated as world-time-individual triples rather than world-time pairs. The upshot is that there can be a variation in truth-value without different propositions being expressed.⁴

²This is Lewis' (1980) terminology. Strictly speaking, Kaplan (1989) talked about circumstances of evaluation, not indices. However, in the interest of preserving a more or less uniform terminology, I am sticking with Lewis' terminology.

³Kaplan (1989) was careful when it came to identifying the content of a sentence in a context with a proposition. As he pointed out, 'This functional notion of the content of a sentence in a context may not, because of the neutrality of the content with respect to time and place, say, exactly correspond to the classical notion of a proposition' (Kaplan, 1989, p. 504).

⁴I am using 'relativism' in a broader sense than MacFarlane (2005b, 2009) insofar as I am not distinguishing between views according to which the relevant parameters are determined by the context of utterance and views according to which they are determined by the context of assessment. Moreover, distinguishing between the context of utterance and the context of assessment also allows us to identify a view according to which the proposition expressed depends on the context of assessment. See e.g. Cappelen (2008a, 2008b) and Weatherson (2009) for relevant

3.2 The Problem of Disagreement

Supporters of relativism have argued that contextualism has problems making sense of disagreement involving predicates of taste.⁵ For instance, take a context in which Mary and John are riding a roller coaster. They both recognise that while John is enjoying the ride, Mary is not, and the conversation in (33) takes place.

- (33) a. John: This is fun.
b. Mary: I disagree, this is not fun.

It is easy to construct similar examples involving ‘tasty’. Imagine that Mary and John are at a party, and that the dialogue in (34) takes place as they are having some chilli.

- (34) a. John: This chilli is tasty.
b. Mary: I disagree, this chilli is not tasty.

In both (33) and (34), it appears correct for Mary to report that she disagrees with John.⁶ However, it is not clear that contextualism respects this judgement. The worry is that if contextualism is correct, and the relevant sentences express different propositions in different contexts, the proposition expressed by the sentence uttered by Mary will not contradict the proposition expressed by the sentence uttered by John. In that sense, Mary will not be denying the proposition expressed by the sentence uttered by John. But then in what sense do they disagree?

discussion. Since I am primarily interested in defending a traditional form of contextualism, these distinctions are not crucial to my overall line of argument. See also Chapter 1.

I also acknowledge that it is somewhat arbitrary to start out by taking indices to be world-time pairs rather than just worlds or world-time-location triples. However, as far as I can see, nothing substantial depends on this.

⁵For a more detailed presentation of the argument against contextualism see e.g. Lasersohn (2005).

⁶The use of ‘I disagree’ may be a bit stilted, but insofar as we want to focus on issues concerning disagreement it makes sense to start out by looking at cases in which one of the parties explicitly reports that she disagrees with the other party. The possible stiltedness of ‘I disagree’ will be relevant in §3.3 and 3.4 when I discuss examples like (35). I will look at other disagreement markers in §3.5.

The problem becomes particularly clear if we consider a rather simple-minded version of contextualism according to which it is always just the tastes of the speaker that are relevant.⁷ In order to appreciate the point it is worth looking at an example which does not involve predicates of taste, but which involves the first-person pronoun ‘I’. The following example is a slightly modified version of an example that Lasersohn (2005) uses to make a similar point:

- (35) a. A: I am hungry.
b. B: I disagree, I am not hungry.

The response in (35a) is clearly inappropriate. It is hard to make sense of the response except as a misunderstanding of what the speaker of (35a) said. If the simple-minded version of contextualism were correct, we might expect (33b) and (34b) to be bad in the same way as (35b), but that does not seem to be the case.

On the other hand, if relativism is correct, the proposition expressed by (33a) and the proposition expressed by (33b) cannot both be true relative to the same index, and in that sense the proposition expressed by (33a) contradicts the proposition expressed by (33b). While more needs to be said before it becomes clear that relativism offers a genuine explanation of the apparent disagreement in (33) and (34), the main question that I am interested in is how problematic examples like (33) and (34) are from a contextualist point of view.⁸ I will argue that these examples are less problematic from a contextualist point of view once we take into account a sufficiently broad range of cases of disagreement.

3.3 More Disagreement Data

As it has been presented here, the argument against contextualism rests on the assumption that Mary must be denying the proposition expressed by the sentence

⁷I am not going to discuss the extent to which contextualists can avoid this problem altogether by adopting a more sophisticated version of contextualism. For attempts to defend contextualism in this way, see e.g. Glanzberg (2007), Recanati (2007, ch. 11), Cappelen and Hawthorne (2009), and Schaffer (forthcoming-a). See e.g. Lasersohn (2005) and MacFarlane (2007) for critical discussion.

⁸For a discussion of how to develop a relativist account of disagreement see e.g. MacFarlane (2007).

uttered by John in order for them disagree. While I am going to argue that we have good reasons to reject this assumption, it might seem plausible if one is only looking at a fairly limited range of cases of disagreement. It is perhaps tempting to think that the examples that we considered in §3.2, should be understood in more or less the same way as the disagreement in examples like (36).

- (36) a. Robert: The Cuban missile crisis took place in 1963.
b. Bobby: I disagree, it took place in 1962, not 1963.

On the assumption that Robert and Bobby are talking about the actual history of the world and are referring to the same historical event, I think it is safe to say that Robert and Bobby disagree about when the Cuban missile crisis took place. By uttering (36b), Bobby is denying the proposition expressed by (36a), the sentence uttered by Robert. More precisely, the proposition expressed by the sentence uttered by Bobby and the proposition expressed by the sentence uttered by Robert cannot both be true.

This is a clear and uncontroversial example of disagreement. Indeed, it might be tempting to think that we can understand all cases of disagreement in dialogues of this form along the same lines as (36), that is, in terms of one of the speakers denying the proposition expressed by the sentence uttered by the other speaker. However, it does not seem that we can fit all cases of disagreement into this mould. I am primarily interested in examples like (32).

- (32) a. Mary: I like this chilli.
b. John: I disagree, it's too hot for me.

We can imagine the conversation in (32) as taking place while Mary and John, as in §3.2, are having some chilli at a party. I take it that John disagrees with Mary when he asserts (32b) in response to (32a) even though the truth of (32a) as uttered by Mary is consistent with the truth of (32b) as uttered by John. If John was denying the proposition expressed by (32a), he would have to think that Mary was somehow insincere or mistaken or confused about her own tastes. Apart from this being odd under normal circumstances, it would be hard to make sense of him saying that the chilli is too hot for him. In other words, we seem to have a case of disagreement in which John is not denying the proposition expressed by the sentence uttered by Mary.

Judgements about examples like (32) are liable to vary to some extent. Even though some find John's use of 'I disagree' a bit stilted, I will assume for now that we want to take the example at face value and say that Mary and John disagree.⁹ In §3.4 I will address some potential concerns about examples like (32), and argue that such examples should be taken seriously.

Examples like (32), if taken at face value, are relevant in the context of evaluating the extent to which the original data present a problem for contextualism. Insofar as it is being assumed that the cases of disagreement that are allegedly problematic from a contextualist point of view have to be understood on the model of (36), examples like (32) show that this cannot be true in general. It is not always the case that disagreement in dialogues of this form is a matter of one of the speakers denying the proposition expressed by the sentence uttered by the other speaker.¹⁰ In that case, we should not necessarily expect that to be true of (33) and (34) either.

- (33) a. John: This is fun.
b. Mary: I disagree, this is not fun.

- (34) a. John: This chilli is tasty.
b. Mary: I disagree, this chilli is not tasty.

⁹Even if we grant that there may be something stilted about the use of 'I disagree' in (32), it is not clear what this shows. The point that the use of 'I disagree' may be stilted does not only apply to examples like (32). These are among the issues that will be addressed in §3.4.

¹⁰Examples like (32) are not the only examples of disagreement that cannot easily be understood along the same lines as (36). Suppose that Pierre and Marie are conducting a series of experiments in order to test a scientific hypothesis that they both believe to be false and that the following conversation takes place after one of the experiments:

Pierre: The hypothesis is false.

Marie: I disagree, we need to do further testing.

In this case we do not need to understand Marie as denying the proposition expressed by the sentence uttered by Pierre. Indeed, since she herself believes that the hypothesis is false, and thus accepts the proposition expressed by the sentence uttered by Pierre, it is more plausible to take her to disagree with Pierre on the grounds that there is insufficient evidence for him to assert that the hypothesis is false without further testing. In other words, Pierre and Marie disagree even though Marie is not denying the proposition expressed by the sentence uttered by Pierre.

Furthermore, examples like (33) and (34) are much less surprising from a contextualist point of view if we take examples like (32) at face value. There are two important points here. The first point to note is that examples like (32) do not lend any kind of support to relativism. Indeed, it is not clear how a plausible account of this kind of disagreement along relativist lines could be constructed. While there may be difficult questions pertaining to the correct semantics for 'likes', there does not seem to be any further need for a relativist or contextualist semantics of the sort that has been proposed for predicates of taste. This means that we have to look elsewhere for an explanation of the disagreement in examples like (32).

The second point is that if we find disagreement when speakers use sentences like 'I like this chilli' to express their personal tastes, it is not surprising, even from a contextualist perspective, that the same is true in the cases involving 'fun' and 'tasty'. It seems reasonable to assume that whatever the right explanation of Mary and John's disagreement in (32) is, that explanation can be extended to cover examples like (33) and (34) as well.

If all of this is correct, it seems to undermine the argument against contextualism. The presence of disagreement in cases involving 'fun' and 'tasty' is to be expected given that we have disagreement in cases involving 'likes'. Moreover, the presence of disagreement in the latter case does not call for a relativist explanation and once this is recognised it is hard to see why it should do so in the former case either. At least the relativists owe us an account of why the cases should be treated differently.

This does not mean that there is nothing to be gained by pursuing other ways of defending contextualism. In particular, it might still be a good idea to improve on the simple-minded version of contextualism according to which it is only the tastes of the speaker which matters. What is relevant may be the tastes of the members of some group that is determined by the context and not just the tastes of the speaker. For instance, it may sometimes be the standards of the wider community that are relevant Recanati (2007, ch. 11). This would make it easier to explain why we take Mary and John to disagree, because we can interpret them as making claims about what is fun and tasty for members of their community. This will become relevant in §3.5.

Still, I think it is important to recognise that contextualists are not necessarily committed to thinking about the disagreement in these terms. Mary and John may disagree even if Mary is not denying the proposition expressed by the sentence uttered by John. This gives contextualists a certain amount of flexibility when it comes to responding to worries about disagreement.

3.4 Doubts about the Data

All of this requires that one takes the data from the previous section at face value. While I am not alone in drawing attention to examples of disagreement like (32),¹¹ some may be inclined to dismiss examples like (32), outright and argue that, contrary to appearances, Mary and John do not disagree after all.

- (32) a. Mary: I like this chilli.
b. John: I disagree, it's too hot for me.

In response to an objection along these lines, it makes sense to compare examples like (32) with examples in which it really does seem wrong to report that the parties disagree. While judgements about the appropriateness of John's response in (32) might be somewhat delicate, it is useful to consider the difference between (32) and (35).

- (35) a. A: I am hungry.
b. B: I disagree, I am not hungry.

As noted earlier, the response in (35b) is clearly inappropriate and it is hard to make sense of it except as a misunderstanding. On the other hand, dialogues like (32) do not seem to exhibit this kind of inappropriateness. We do not have the same problem making sense of John's response in (32). I take this to count against someone who wants to say that there is no disagreement in examples like (32). At the very least, one would have to explain why John's response in (32) is more natural and appropriate than the response in (35).

In §3.3 I acknowledged that some find the use of 'I disagree' in (32) a bit stilted. But to the extent that I am willing to grant that, I also think that to some extent can be said of (33) and (34) as well.

¹¹See e.g. Weatherson (2009) who points out that this sort of disagreement is known from debates about non-cognitivism in moral philosophy.

- (33) a. John: This is fun.
b. Mary: I disagree, this is not fun.
- (34) a. John: This chilli is tasty.
b. Mary: I disagree, this chilli is not tasty.

It is not clear that there are sufficient grounds for dismissing examples like (32) while continuing to treat examples like (33) and (34) as evidence against contextualism. In general, it is problematic to be dismissive of examples like (32) if one wants to make positive use of other disagreement data. There might very well be reasons to conclude that examples like (32) are not cases of genuine or real disagreement in some interesting sense, but then we should be prepared to say the same thing about other putative cases of disagreement as well. In this context it is also worth noting that some philosophers have expressed doubts about the judgements about disagreement that relativists rely on.¹² While I am not advocating this line of response on behalf of contextualism, it makes it somewhat awkward from a dialectical point of view for relativists to be dismissive of other putative cases of disagreement. In the absence of a more significant difference between examples like (32) and examples like (33) and (34), it is not clear that relativists are warranted in dismissing examples like (32) while treating examples like (33) and (34) as unproblematic.

Even if one does not want to outright dismiss examples like (32), one might think that since the denial in (32b) cannot plausibly be interpreted as targeting the proposition expressed by (32a) it could target some other proposition conveyed by Mary's utterance. Perhaps a relativist could argue that Mary's utterance conveys the proposition that the chilli is tasty, with this being understood in relativist terms as having different truth-values relative to different individuals. Insofar as relativists are in a position to explain the disagreement in examples like (33) and (34), they could extend that explanation to examples like (32). On this picture, relativism would be needed to make sense of the disagreement after all.

While this is an interesting suggestion that deserves further investigation, there are some questions that need to be addressed. For one thing it is not obvious why we should take Mary to convey this particular proposition by her

¹²See e.g. Cappelen and Hawthorne (2009, ch. 4) and Stanley (2005, p. 144).

utterance. It is tempting to think of this in terms of a conversational implicature, but how does this implicature arise?¹³ A natural suggestion is to appeal to considerations involving relevance.¹⁴ While there might be other explanations available, this looks like the most obvious strategy. The basic idea is that what is relevant is whether the chilli is tasty and not just whether Mary likes, it and that is why she is understood as conveying that the chilli is tasty.

However, given the context in which Mary utters (32a), it is not clear why her contribution to the conversation should not count as sufficiently relevant unless she is understood as conveying something beyond that she likes the chilli, and more specifically, that the chilli is tasty. If they had been talking about what kind of food they find tasty, that might be a reasonable assumption, but that does not have to be the case. They could just as well have been talking about what they like about the party or what is happening there. Moreover, this does not seem to affect the judgement that Mary and John disagree in (32). This is not to rule out an explanation in terms of relevance, but only to say that the availability of such an explanation cannot be taken for granted.

Another question is whether the presence of a conversational implicature would be sufficient to explain the disagreement data. It is not clear whether it

¹³An anonymous referee for the *Australasian Journal of Philosophy* pointed out to me that the dialogue in (32) gets worse if we replace 'it's too hot for me' with 'I am tired', even though the latter may also be a good reason for not enjoying the chilli. A possible explanation is that the taste of the chilli is somehow at stake, and being tired does not bear on that. However, even if it is true that the taste of the chilli is somehow at stake, it is far from clear that this shows that an account of the sort that would favour the relativists is correct.

The same referee suggested that another possible relativist response is to say that (32a) entails that Mary finds the chilli tasty and that allows the denial to target the proposition that the chilli is tasty. I will continue to focus on implicature-based proposals, but it is worth noting that this is might be an interesting alternative. However, more work remains to be done before it is clear that the suggestion is plausible, and I will not pursue it here.

¹⁴If we take Grice's (1989a) framework as a starting point, it seems natural to appeal to the maxim of Relation. As far as the maxim of Quality and the maxim of Manner are concerned, Mary's contribution to the conversation is presumably both true and based on good evidence, and it is not obscure or ambiguous. Matters are more complicated in the case of the maxim of Quantity. However, to the extent that it makes more sense to appeal to the maxim of Quantity, this strategy seems to face the same questions as a strategy based on the maxim of Relation. In particular, it is not clear why Mary's contribution to the conversation should not count as sufficiently informative even if she is only conveying that she likes the chilli.

is always appropriate to use ‘I disagree’ when rejecting an implicature. Let us consider a case inspired by one of Grice’s (1989a, p. 37) original examples. A and B are private detectives hired by Mrs. X to spy on her husband, Mr. X, and they are talking about what he is up to.

- (37) a. A: Mr. X is meeting a woman tonight.
b. B: I disagree, he is meeting his wife.

Following Grice, I take it that A is implicating that the woman Mr. X is meeting is not his wife, and B is naturally understood as rejecting that implicature. Still, there seems to be something awkward about B’s use of ‘I disagree’ in (37b). These judgements are admittedly fairly subtle, and I am reluctant to put too much weight on a single example, but the point is to raise a question about whether the use of ‘I disagree’ is always appropriate when rejecting an implicature in this way. If it is not, there is at least some additional work to be done to explain why it is appropriate in the case of (32). If one is not impressed by examples like (32), this may not be a particularly pressing worry, but insofar as one finds John’s response in (32) appropriate and the response in (37) inappropriate, this is something that deserves further investigation.¹⁵

These considerations are not meant to rule out the possibility of saying that what is going on in (32) is that John is rejecting an implicature. The point is that relativists are not entitled to assume that the disagreement can be explained in this way. There are questions that need to be answered and it remains to be seen whether there is a satisfactory story to be told.

3.5 Other Disagreement Markers

Even if one accepts a lot of what I have said so far, there is still a worry that needs to be addressed. The worry is that I have been focusing on the wrong sort of data, and that the relevant data involve a more restricted notion of disagreement. The cases which I have been concerned with involve the use of ‘I disagree’,

¹⁵This does not amount to the claim that the use of ‘I disagree’ is never appropriate when targeting an implicature. Given that there are significant differences between different kinds of implicatures, it would perhaps not be surprising to find that the use of ‘I disagree’ is more or less appropriate in cases involving other kinds of implicatures. However, that would only illustrate the need to get clear on the details of the account.

but there could be other disagreement markers in English that are more suited to the argumentative purposes of the relativists by allowing us to focus on a more restricted notion of disagreement. While relativists do not have to claim that the relevant notion of disagreement is always tied to the use of a specific disagreement marker, it would be useful for them if they had a way of showing that our judgements are tracking the relevant notion of disagreement. If it turns out that our judgements about cases involving other disagreement markers differ from our judgements about cases involving ‘I disagree’, there might be a concern that it will be difficult to extend the contextualist response that I have proposed to these cases. I will look at two such proposals and I will argue that neither of them provides relativists with the desired results.

One option is to argue that the relevant cases involve the use of ‘that’s not true’ or ‘that’s false’. Insofar as we want to focus on cases of disagreement in which what is being denied is the proposition expressed by the sentence that is uttered, this looks like a plausible candidate. Furthermore, this appears promising from a relativist perspective given that John’s response in (38) is clearly inappropriate.

- (38) a. Mary: I like this chilli.
b. John: That’s not true, it’s too hot for me.

In this respect (38) differs from (32).

- (32) a. Mary: I like this chilli.
b. John: I disagree, it’s too hot for me.

This does not seem to get the relativist very far though. The problem with this approach is that there is also something unnatural about (39) and (40).

- (39) a. John: This is fun.
b. Mary: That’s not true, this is not fun.

- (40) a. John: This chilli is tasty.
b. Mary: That’s not true, this chilli is not tasty.

For instance, we can consider the dialogue in (39) as taking place in the same context as before. John is enjoying the ride, whereas Mary is scared and sick and

clearly not enjoying it. In this context, there seems to be something unnatural about Mary's response in (39). Similarly, if Mary and John are just trying the chilli at a party, there is something awkward about Mary's response in (40). In this respect (39) and (40) differ from (33) and (34).

- (33) a. John: This is fun.
b. Mary: I disagree, this is not fun.

- (34) a. John: This chilli is tasty.
b. Mary: I disagree, this chilli is not tasty.

But (39) and (40) are examples of the sort that were meant to present a problem for contextualism. If these dialogues do not sound natural, I assume that this should not be taken as evidence against contextualism insofar as this is just what we would expect if contextualism were true. Judgements may vary somewhat with respect to these examples and the extent to which the responses in (39) and (40) are inappropriate, but I do not take these examples to be very impressive as evidence for relativism, and that is what is important for our present purposes.

While it needs to be acknowledged that (38) sounds worse than (39) and (40), that is not necessarily unexpected from the point of view of a reasonably flexible and sophisticated version of contextualism. For instance, in the case of (39), we can try to make sense of Mary's behaviour by taking her to be interpreting John as making a stronger claim about what is fun for both of them even if that is not natural given the context. However, in the case of (38), Mary makes it explicit that she is talking about what she likes and that means that such a stronger reading is unavailable.

There are cases involving predicates of taste in which a response of this sort is more natural.¹⁶ If Mary and John are opening a restaurant in a reasonably civilised location, and they are talking about what they are going to serve in the restaurant, it seems more natural for Mary to respond to John's utterance of (41a) with (41b).

- (41) a. John: Rotten shark is tasty.
b. Mary: That's not true, rotten shark is not tasty.

¹⁶Thanks to Jonathan Schaffer for making me appreciate this point.

While Mary's response in (41b) sounds more natural than the response in (40b), this is not a problem for contextualism. Given a very reasonable contextualist story, this is just what we should expect. In this context, it is reasonable to think that is not just the tastes of the speaker which are relevant, but the tastes of the members of a larger community of which Mary is also a member. In this case the contextualists can simply understand Mary as denying the proposition expressed by the sentence uttered by John. Moreover, it is not entirely clear what a relativist explanation of the difference between this example and examples like (39) and (40) would look like. Given the relativist position we have been considering so far, there does not seem to be too much of a difference between them. In other words, there is no threat to contextualism from examples like (41), and the examples which would pose a threat to contextualism, examples like (39) and (40), do not sound very convincing in the first place. If anything, our judgements about these cases seem to fit well with what we would expect if contextualism were true.

Having said that, there are other disagreement markers that are worth examining. Stephenson (2007) is for instance careful to point out that she is focusing on a more restricted notion of disagreement that is tied to the use of expressions like 'nuh-uh' and 'no', and she allows for the possibility that whether two speakers disagree or not is a broader phenomenon. Kölbel (2002, 2004) and Lasersohn (2005) also focus on examples of this sort. Initially this seems more promising, since there is no problem about either (42) or (43).

- (42) a. John: This is fun.
 b. Mary: Nuh-uh/No, this is not fun.
- (43) a. John: This chilli is tasty.
 b. Mary: Nuh-uh/No, this chilli is not tasty.

However, it is not clear that focusing on disagreement markers like 'nuh-uh' or 'no' is going to be enough for the relativists. Consider the dialogue in (44).

- (44) a. Mary: I like this chilli.
 b. John: Nuh-uh/No, it's too hot for me.

Again we can again imagine that the conversation in (44) takes place while Mary and John are having some chilli at a party. Is John's response appropriate in this

context? While some may be uneasy about this example, the response does not seem to be altogether unnatural in the relevant context. Moreover, as I stressed earlier, I think it is important to compare the relevant examples with examples such as (45).

- (45) a. A: I am hungry.
b. B: Nuh-uh/No, I am not hungry.

I take the dialogue in (44) to be more natural than the one in (45), even though one might expect (44) to be just as bad as (45). This is at least something which requires explanation and there does not seem to be any reasons to think that this explanation would ultimately give us any reason for preferring relativism over contextualism, for reasons that we have already discussed.

If one is still not convinced about examples like (44), that does not mean that the use of disagreement markers like ‘no’ and ‘nuh-uh’ in examples like (42) or (43) show that Mary must be denying the proposition expressed by the sentence uttered by John. There are other examples that one might want to take into account. For instance, consider an example like (46).

- (46) a. Mary: I recommend this chilli.
b. John: Nuh-uh/No, it’s too hot for me.

John’s response in (46) seems to be appropriate. However, it is implausible to interpret him as denying the proposition expressed by the sentence uttered by Mary. Moreover, it is plausible to think of ‘fun’ and ‘tasty’ as often being used to recommend something. This suggests that it would be a mistake to focus exclusively on examples like (44).

In the end, if I am far from convinced that the use of ‘nuh-uh’ or ‘no’ as a disagreement marker instead of ‘I disagree’ makes enough of a difference for the relativist to be able to mount an effective argument against contextualism. There are bound to be issues regarding when it is preferable to use ‘I disagree’ and when it is preferable to use ‘no’ and ‘nuh-uh’, but in the absence of more clear-cut data I would be wary of investing too much in this distinction. On the overall assessment of the disagreement data considered in this section, I do not think that contextualism fares too badly. When it is less natural to use a disagreement marker in examples like (32), as in the case of ‘that’s not true’, it also less

natural in examples like (33) and (34). However, when it is more natural to use a disagreement marker in examples like (33) and (34), as in the case of ‘nuh-uh’ and ‘no’, that is also true of examples like (32). While I do not pretend that these matters are in any way clear-cut, I do not think that this amounts to anything like a good case for preferring relativism over contextualism.

3.6 Conflicting Attitudes

My main concern in this paper is to argue that there are other cases of disagreement that we ought to take into account when we are evaluating the claim that considerations involving disagreement give us a reason for preferring relativism over contextualism. Examples like (33) do not call for a relativist treatment.

- (32) a. Mary: I like this chilli.
b. John: I disagree, it’s too hot for me.

Furthermore, once we take examples like (32) into account, examples like (33) and (34) are less problematic from a contextualist point of view.

- (33) a. John: This is fun.
b. Mary: I disagree, this is not fun.
- (34) a. John: This chilli is tasty.
b. Mary: I disagree, this chilli is not tasty.

While this does not presuppose any specific explanation of the data, I want to offer some speculations as to what an explanation that is compatible with a contextualist treatment of predicates of taste might look like. The idea I am interested in is to view disagreement as a matter of the parties having incompatible or conflicting attitudes. Two parties disagree just in case there is something that they have conflicting attitudes towards. This sometimes means that there is a proposition such that one party believes that proposition and the other party believes its negation, but that does not always have to be the case. Just as two parties may have conflicting beliefs, they may also have conflicting desires or preferences.

A detailed examination of the relevant issues is beyond the scope of this chapter, and the goal is only to present the basic idea.¹⁷ Still, if these speculations turn

¹⁷See Chapter 2 for further discussion.

out to be on the right track, it would lend further support to the main point by showing that there is a way of thinking about the disagreement in the relevant examples that does not favour relativism, or any other semantic theory for that matter.

To the extent that disagreement may involve different sorts of attitudes there is a sense in which there are varieties of disagreement. It would therefore be a mistake to think that all disagreement should be understood in exactly the same way as examples like (36), in which there is a proposition such that one party believes that proposition and the other party believes its negation.

- (36) a. Robert: The Cuban missile crisis took place in 1963.
b. Bobby: I disagree, it took place in 1962, not 1963.

Furthermore, having a more liberal view of disagreement makes it easier to make sense of puzzling examples like (32). If this suggestion is on the right track, we can say that Mary and John disagree in virtue of having conflicting attitudes towards the chilli, and not in virtue of John denying the proposition expressed by the sentence uttered by Mary. We can also approach Mary and John's disagreement in (39) and (40) in much the same way since they presumably have conflicting attitudes towards the roller coaster ride and the chilli in those cases as well.¹⁸ On this picture, it is the attitudes of the parties that matter, and this way of thinking of the disagreement is to a large extent independent of the semantic issues. In particular, it does not matter whether we adopt a contextualist or relativist semantics for predicates of taste.

There is something *prima facie* plausible about the thought that the disagreement in (32) has something to do with Mary liking the chilli and John disliking the chilli. Moreover, this way of thinking about disagreement is not entirely without precedent. It is similar to ideas usually associated with expressivist theories of the sort that figure prominently in debates in moral philosophy.¹⁹ Roughly

¹⁸The idea that the disagreement in cases involving predicates of taste can be understood in terms of incompatible attitudes towards something that is not a proposition, has also been suggested by Maudlin (2007).

¹⁹See e.g. Stevenson (1937, 1944, 1963) who distinguishes between what he calls 'disagreement in belief' and what he calls 'disagreement in attitude'. The way that I think about examples like (32) is very closely related to his notion of disagreement in attitude. See e.g. also Blackburn (1984, 1998).

speaking, expressivists take certain expressions to express attitudes, say, approval or disapproval, praise or resentment. This can be contrasted with the sort of truth-conditional semantics that I have been presupposing throughout this discussion, according to which the content of an expression is represented by a function from indices to extensions.²⁰ Since I am mainly interested in the debate between contextualists and relativists, I do not want to enter into a discussion about whether expressivism could provide a viable alternative to a contextualist or relativist semantics for predicates of taste. However, it is interesting to note that expressivists have claimed to be in a good position to make sense of disagreement by taking it to involve a conflict of attitudes in more or less the same way that I am suggesting.²¹

Having said that, it is important to recognise that thinking about disagreement in this way still does not force us to adopt a particular semantic theory. One can think about disagreement in this way without endorsing expressivism. In order for the speakers to disagree, they need to have certain attitudes, but there is no requirement that the relevant expressions express these attitudes in the sense that expressivists are interested in.²² In the case of (32), Mary is reporting that she likes the chilli, and it is also reasonable to suppose that a sincere utterance of (33a) or (34b) is typically, though not invariably, accompanied by the speaker having a certain attitude towards the roller coaster ride or the chilli.

Even if this suggestion enjoys some *prima facie* plausibility and has some precedent in the literature on expressivism, there are many outstanding questions. Apart from clarifying the notion of ‘conflicting attitudes’, there is also more work to be done when it comes to relating what is going on at the level

²⁰Blackburn and Gibbard are often mentioned as prominent contemporary defenders of expressivism in moral philosophy. See e.g. Blackburn (1984, 1998) and Gibbard (1990, 2003). However, I will not delve into the subtle issues about how to interpret their views.

²¹Commenting on a case of moral disagreement over contraception, Blackburn claims that the expressivist theory ‘locates the disagreement where it should be, in the clash of attitudes towards contraception’ (Blackburn, 1984, p. 168). See also e.g. Blackburn (1998, p. 69).

²²To the extent that expressivists are in a position to tell a plausible story about disagreement in terms of conflict attitudes, it is not clear why a contextualist, or a relativist for that matter, cannot tell more or less the same story. Dreier (1999) makes a similar point with respect to contextualism and expressivism in moral philosophy. He argues that a contextualist about moral judgements can follow the expressivist when it comes to explaining moral disagreement, for instance by talking about disagreement in attitudes or norms. See e.g. also Jackson and Pettit (1998).

of mental states to what is going at the level of language and discourse.²³ For instance, it would be interesting to see how this way of looking at disagreement can be integrated into a more general account of agreement and disagreement in discourse.²⁴

Conclusion

In this chapter I have sought to answer a specific challenge to contextualist treatments of predicates of taste, namely that such treatments leave us unable to account for certain cases of disagreement. It has been argued that the disagreement data favour a relativist account of predicates of taste. Contrary to this line of argument, I have attempted to show that considerations involving disagreement ultimately do not provide evidence for relativism. Once we look at a broader range of cases, the original cases of disagreement that were meant to lend support to relativism seem much less surprising from a contextualist perspective. While this does not show that a relativist treatment of predicates of taste is incorrect, I take it to be significant in light of the emphasis that relativists have placed on considerations involving disagreement.

No doubt there is much more that can be said about these matters, but I offer these considerations as a challenge to those who think that a relativist treatment of predicates of taste is to be preferred over a contextualist treatment for reasons having to do with disagreement. I also hope that what I have said suggests further avenues that can be explored in this debate.

²³An anonymous referee for the *Australasian Journal of Philosophy* suggested an example which illustrates some of the difficulties concerning when it is legitimate to appeal to a conflict of attitudes when characterising a case of disagreement. Suppose that it is possible to like or dislike something without being aware of it, and that there is a machine that can detect these attitudes. If the dialogue in (32) took place solely on the basis of Mary and John checking the machine, it seems that John's response would be less natural. If that is right, there must be a difference between the cases that gives rise to different judgements about disagreement.

²⁴See e.g. Geurts (1998), Maier and van Der Sandt (2003), and Asher and Lascarides (2009).

Chapter 4

Epistemic Modals

Introduction

In this paper I will be defending traditional contextualist treatments of epistemic modals against certain objections that have recently been raised against this view. According to contextualism, epistemic modals are context-dependent in the familiar sense that they express different propositions in different contexts. While a contextualist treatment of epistemic modals is often regarded as at least having some initial attraction, it has been argued that such a view yields incorrect predictions about cases involving disagreement and retraction. If we assume that contextualism is correct, we are supposedly unable to explain apparently mundane cases of disagreement involving epistemic modals. Furthermore, it is argued that contextualism predicts that it would be wrong for speakers to retract their earlier claims in situations in which such retractions seem to be perfectly natural. Several critics of contextualism have therefore suggested an alternative semantics for epistemic modals. In particular, several recent critics of contextualism have advocated some kind of relativist treatment of epistemic modals.¹

¹See e.g. Egan et al. (2005), Egan (2007), Stephenson (2007) and MacFarlane (forthcoming-a) who have defended a relativist treatment of epistemic modals. However, it should be noted that there are some differences between the positions that they defend. For various attempts to defend a more traditional contextualist approach, see e.g. Papafragou (2006), Hawthorne (2007), Wright (2007), Schaffer (forthcoming-a), and von Fintel and Gillies (2008, forthcoming).

I want to argue that these problems have been exaggerated, and that they do not force us to abandon the familiar contextualist treatment epistemic modals. There might be other reasons for giving up the simple contextualist picture, but alleged problems having to do with disagreement and retraction do not provide us with reasons to look for a different semantics. In order for these considerations to provide an argument against contextualism, the disagreement and retraction in question would have to target the proposition expressed by the relevant sentences. However, there are reasons for being suspicious of this assumption. For instance, as von Fintel and Gillies (2007, 2008) have observed, there are important similarities between epistemic modals and first person uses of attitude verbs like ‘thinks’ when it comes to disagreement data. The general worry is that the disagreement and the retraction are targeting the embedded proposition. In that case, the data does not tell us much about the semantics of epistemic modals. My goal is to pursue this as a general line of response on behalf of contextualism and to evaluate its general merits and shortcomings.

In §4.1 I outline the difference between contextualist and relativist treatments of epistemic modals. In §4.2 I present the disagreement and retraction data that allegedly present a problem for contextualist treatments of epistemic modals. In §4.3 I present some further data that are meant to show that these matters are also complicated from a relativist point of view. In §4.4 I briefly discuss the possibility that the relevant patterns of disagreement and retraction can be the result of a more general psychological bias. In §4.5 I propose that the disagreement and retraction should be seen as targeting the embedded proposition and not the proposition expressed by the relevant sentence. In §4.6 I discuss an observation made by Stephenson (2007), and argue that it does not present any additional problems for contextualism. In §4.7 I argue that it is not a plausible constraint on disagreement that disagreement can only target a proposition that is believed or asserted. In §4.8 I also reject the corresponding constraint for retraction. In §4.9 I suggest that we can understand the relevant cases of disagreement involving epistemic modals in terms of the parties having conflicting credences rather than conflicting beliefs. In §4.10 I discuss the view that sentences containing epistemic modals may be used to perform two distinct speech acts, and suggest that this can be a way of getting a better grip on the retraction data. In §4.11 I briefly address problems involving agreement and some further

problems involving probability operators like ‘probably’.

4.1 Contextualism and Relativism

In this chapter, I will be focusing on issues having to do with epistemic uses of modal expressions like ‘might’ and ‘must’. For the most part, the discussion will focus on possibility modals like ‘might’. In this section, I will briefly sketch how epistemic modals like ‘might’ can be understood along contextualist and relativist lines within a fairly standard framework in which modals are treated as quantifiers over possible worlds, setting the stage for the following discussion.

What does it mean to say that the relevant modality is epistemic? Loosely speaking, we can think of epistemic modality as somehow being connected with what is compatible with a certain body of knowledge or evidence. According to a familiar account of modals that I will be working with here, modals can be thought of as quantifiers over possible worlds.² In the case of epistemic modals, the set of possible worlds we are quantifying over are those which are compatible with the available information. It will be convenient to identify the available information with what is known, and I will do so for the purpose of the following discussion. Simplifying somewhat, we can then construe an epistemic use of ‘might’ as involving existential quantification over the possible worlds compatible with the available information. The idea behind this can be illustrated by means of an example. Take (47) as involving an epistemic use of ‘might’.

(47) It might be raining.

Simply put, the idea is that (47) is true in a possible world w just in case there is some possible world w' which is compatible with the available information in w in which it is raining. But what is the relevant information?

According to the standard contextualist treatment of epistemic modals, the relevant information is determined by the context of use. Sentences like (47) express different propositions in different contexts, depending on what the relevant information is.

²For some important work on modality within the possible worlds framework, see Kratzer (1977, 1981, 1991b). See Portner (2009) for an introduction to this kind of framework and linguistic theories of modality in general.

A very simple and straightforward version of contextualism would be the view that the information is to be identified with what is known by the speaker at the time of the utterance. Suppose that Sally utters (47). In that case, (47) is true just in case there is a possible world compatible with what she knows in which it is raining. However, there are other ways of implementing the basic idea behind the contextualist proposal. One could for instance identify the relevant information with the combined knowledge of the conversational participants. That would make the truth of (47) dependent on what is known by the conversational participants in general and not just the speaker. The details do not matter at this point. A different choice point is whether one wants to say that it is what is known or what the speaker or the conversational participants could come to know through some further, but presumably limited, investigation.³ For the sake of simplicity I will be assuming that we are dealing with a simple version of contextualism unless it is explicitly stated otherwise. This will allow me to focus on the issues that I want to focus on, but it does not mean that I endorse this version of contextualism. What matters is that contextualism makes epistemic modals context-dependent in the familiar sense that sentences that contain epistemic modals express different propositions in different contexts.

Relativism has emerged as an important competitor to contextualism in recent debates about epistemic modals. Relativist treatments of epistemic modals differ from contextualist treatments insofar as they do not take sentences like (47) to express different proposition in different contexts. Instead they express propositions that are true or false relative to an appropriate parameter. For the purpose of this discussion we can think of the relevant propositions as having different truth-values relative to different information states. According to the version of relativism proposed by MacFarlane (forthcoming-a), the relevant information is not determined by the context of use, but by the context of assessment.⁴ Roughly speaking, a context of use is understood as a context in a which

³See e.g. Hacking (1967) and DeRose (1991) for discussion about these and related matters.

⁴If we accept this way of construing the debate, the difference between contextualism and relativism comes down to two things. First, it comes down to whether sentences containing epistemic modals express different propositions in different contexts or whether they express propositions that vary in truth-value. Second, it comes down to whether the relevant information is determined by the context of use or the context of assessment. These things can come apart and that means that there other possible views that one might want to explore. For the purpose of the following

a sentence is used, whereas a context of assessment is understood as a context in which the use of a sentence is being assessed. In other words, epistemic modals are context-dependent, but in the less familiar sense that they are sensitive to features of the context of assessment. MacFarlane and others have argued that this allows us to make sense of certain cases of disagreement and retraction involving epistemic modals which contextualists are supposedly unable to explain. In the following discussion, it will be convenient to focus on relativism as an alternative to contextualism. However, my main concern is not going to be relativism as such, or indeed any other alternative to contextualism, but the general problems involving disagreement and retraction facing a contextualist treatment of epistemic modals.

4.2 The Data

While I am open to the possibility that there are all sorts of problems facing a simple contextualist account of epistemic modals, I am going to focus exclusively on issues having to do with disagreement and retraction.⁵ One worry concerning contextualist treatments of epistemic modals is that they cannot handle cases of disagreement in a satisfactory manner. Such worries have been raised by, among others, Stephenson (2007) and MacFarlane (forthcoming-a). Even simple cases of disagreement seem to be problematic from a contextualist point of view. In order to illustrate this point, von Fintel and Gillies (2008) use an example from Kratzer (1991a).

Suppose a man is approaching both of us. You are standing over there. I am further away. I can only see the bare outlines of the man. In view of *my* evidence, the person approaching might be Fred. You know better. In view of *your* evidence, it cannot possibly be Fred, it must be Martin. If this is so, *my* utterance of [(48)] and *your* utterance of [(49)] will both be true.

(48) The person approaching might be Fred.

discussion, I will be ignoring these possibilities. See also Chapter 1.

⁵For other problems facing a simple contextualist treatments of epistemic modals, see e.g. Yalcin (2007, forthcoming).

(49) The person approaching cannot be Fred.

Had *I* uttered [(49)] and *you* [(48)], both our utterances would be false. (Kratzer, 1991a, p. 654, original emphasis)

What von Stechow and Gillies observe, is that there appears to be a sense in which the speaker of (48) and the speaker of (49) disagree. However, it is not clear that this is something we predict if we adopt a very simple version of contextualism according to which the relevant information is what is known by the speaker. In that case, (48) would be true as uttered by the first speaker just in case it is compatible with what the first speaker knows that the person is Fred. On the other hand, (49) would be true as uttered by the second speaker just in case it is incompatible with what second speaker knows that the person is Fred. But then (48) and (49) would express compatible propositions. There does not appear to be a proposition such that the first speaker believes that proposition and the second speaker believes its negation.

In order to elicit judgements about disagreement, it is useful to look at what happens in a dialogue between speakers who have access to different information. These kinds of cases have played a central role in the debate, and they will play an important role in my discussion as well. Suppose that Mary and John are trying to find out where Harry is and that they have so far been unable to find him. John does not know that Mary has just been looking for Harry in the office and that he was not there.

- (50) a. John: Harry might be in the office.
b. Mary: No, he isn't. I just checked and he wasn't there.

Mary seems to disagree with John in the dialogue above. Indeed, her response seems perfectly natural. But if we assume a contextualist treatment of 'might' according to which the relevant information is what is known by the speaker, it is not clear that we can make sense of this. In that case, it would be wrong for Mary to deny John's claim. If the contextualist treatment is correct, (50a) is true as uttered by John since it is compatible with what he knows that Harry is in the office. Of course, it is not compatible with what Mary knows that Harry is in the office, but this does not solve the problem as far as the contextualists are

concerned, since they take that to be perfectly compatible with the truth of what John said.

This is not to say that we never disagree with and correct other people regarding their own mental states. Sometimes it may be reasonable to assume that someone is mistaken in that way. However, in this case there does not seem to be any reason to think that this is what Mary is doing. Why would she think that John is mistaken about what he knows? We can even imagine that Mary knows that John is not aware of her having looked for Harry in the office. In that case, it would be extremely odd for her to deny that it is compatible with what John knows that Harry is in the office, but her response still seems to be fine. For the sake of simplicity, I am simply going to assume throughout the following discussion that the speakers are not being insincere or mistaken about their mental states.

If we instead adopt a relativist semantics for epistemic modals, we seem to be in a better position to explain what is going on. In that case, the truth-value of (50a) depends on the information available in the context of assessment. Assuming that we are dealing with a fairly simple version of relativism, it is what the assessor knows that is relevant. In Mary's context of assessment, (50a) is false, since it is not compatible with what she knows that Harry is in the office. Therefore, her response is appropriate. She is right to deny (50a), since it is false as assessed by her. We are then supposed to conclude that cases like (50) provide support for a relativist treatment of epistemic modals.

In addition to the cases involving disagreement or correction presented above, cases of retraction are also cited as a potential problem for contextualism. MacFarlane (forthcoming-a) makes important use of such cases of retraction when arguing against contextualism and in favour of relativism.⁶

- (51) a. John: Harry might be in the office.
b. Mary: No, he isn't. I just checked and he wasn't there.
c. John: Okay, then I guess I was wrong.

⁶Note that Stephenson (2007) has a different view of the retraction cases. She does not appeal to retraction data in the way that MacFarlane does. In fact, she treats the retraction data as presenting a potential problem for her view. I am not going to discuss her views on retraction here.

As MacFarlane and others have pointed out, it seems quite natural for John to retract his original claim when confronted with Mary's response. In addition, the opponents of contextualism also tend to emphasise that it appears to be strange for him to stand by his original claim and refuse to retract it. Compare John's response in (51c), which sounds natural, with his response in (52c), which is not very natural.

- (52) a. John: Harry might be in the office.
b. Mary: No, he isn't. I just checked and he wasn't there.
c. John: Okay, then he can't be in his office. But I still stand by what I said a second ago.

Contextualism seems to be ill-suited to explain this data. If contextualism were correct, the proposition expressed by (51a) would be true and there would apparently be no reason for John to retract it in the face of Mary's response. Furthermore, it is not clear why John could not stand by his original claim. Contextualism seems to yield the incorrect prediction that the response in (52c) should be fine whereas it is in fact at best stilted and unnatural.

Again, proponents of relativism claim to be in a position to offer an explanation of what is going on. Since John's epistemic position has improved, it is no longer compatible with what he knows that Harry is in the office. In the new context of assessment, the original claim is false, and hence it makes sense for John to retract it. Thus it seems that relativism is better suited than contextualism to explain the retraction data.

4.3 More Data

There is reason to think that what was said in the previous section is a bit too quick as a characterisation of the relevant data. The picture is more complicated, at least as far as the retraction data is concerned. As von Stechow and Gillies (2008) point out, it sometimes sounds appropriate for the speaker to stand by her initial claim.⁷ For instance, Alex's response in (53c) sounds quite appropriate and natural.

- (53) a. Alex: They keys might be in the drawer.

⁷See e.g. Hacking (1967) for a similar observation.

- b. Billy: [Looks in the drawer, agitated.] They're not. Why did you say that?
- c. Alex: Look, I didn't say they *were* in the drawer, I said they *might be* there—and they might have been. Sheesh.

This is not very surprising from a contextualist point of view. On this view, (53a) is predicted to be true as uttered by Alex and it makes sense that she is in a position to defend her assertion.⁸ Matters are somewhat more complicated if we look at cases like this from a relativist perspective. Once Alex learns that the keys are not in the drawer, (53a) is predicted to be false relative to her context of assessment. It is therefore unclear how she could continue to defend her assertion. However, one might wonder whether Alex is really defending the truth of her assertion or only its appropriateness given what she knew at the time. If it is only the latter, relativists might still be able to make sense of what is going on. For one thing, it is not clear that the second occurrence of 'might' in (53c) is epistemic. This gives relativists some leeway when dealing with cases like this.⁹ What matters for our purposes is that an explicit retraction is not always called for. Indeed, it looks like the speaker has a choice in the cases under discussion. She can either retract or she can stand by her original claim. One way or another, this is presumably something that we want a plausible account of the data to capture.

Interestingly, there are also cases involving epistemic modals in which it sounds less natural for a speaker to retract her original claim. For instance, consider the following dialogue between Mary and John:

- (54) a. John: Joe might be coming to the party, but it would be very surprising if he did.
- b. Mary: Actually he won't be coming to the party. I just talked to him a minute ago.
- c. John: Okay, then I guess I was wrong.

There is nothing odd about John's original utterance in (54a), but the retraction in (54c) is somewhat odd. At the very least, it comes across as being extremely concessive. Insofar as the retraction is inappropriate, that is surprising from a

⁸As Josh Dever (p.c.) pointed out to me, there are still difficult problems, even for a contextualist, when it comes to making sense of the interaction between tense and modals.

⁹See e.g. MacFarlane (forthcoming-a) for relevant discussion.

relativist point of view. At least at first sight, (54a) appears to be a conjunction in which the first conjunct contains an epistemic modal. After John learns that Joe is not coming to the party, the first conjunct is predicted to be false relative to his context of assessment. But then we would expect it to be appropriate for him to retract his original claim. No doubt there are things to be said on behalf of relativism at this point, but the take home message is that the data is more complicated than one might initially have supposed, and it is not even clear that the relativists get all the cases right. A successful account of the retraction data would at least have to account for the appropriateness of John's retraction in (51), the appropriateness of Alex defending her original assertion in (53), and possibly even the inappropriateness of John's retraction in (54).

Even if one is not particularly convinced by the relativist treatment of the data, there is at least something that calls for an explanation here if one is attracted to the fairly simple contextualist picture that I sketched earlier. If anything, it looks like the data is more complicated and harder to explain than one might have thought. In the following I will discuss and defend a possible answer on behalf of contextualism.

4.4 Psychological Bias

Before I turn to the main discussion of the disagreement and retraction data, there is another worry concerning the data that I want to mention. The worry is that some sort of general psychological bias could be at work in the cases under discussion. If that is the case, it is not clear whether the data should be taken to provide evidence against contextualism. After all, we might not want the presence of a general psychological bias to be reflected in our semantics for epistemic modals. In recent work on knowledge ascription, Nagel (2010) has drawn attention to a phenomenon known as 'epistemic egocentrism'. This is a robust psychological bias that affects our ability to accurately represent the epistemic perspectives of less informed subjects. Apparently, we have a hard time suppressing information that is available to us, but not to the subjects we are thinking about, when evaluating their judgements. This includes, but is not restricted to, situations in which we evaluate our own past mental states.

The broader problem of epistemic egocentrism affects not only our efforts to reconstruct our own past mental states, but also our efforts to judge the mental states of less informed others: we overestimate the extent to which they share our beliefs, attitudes and concerns, even in the face of feedback to the contrary, and are surprisingly unaware of the extent to which we do this. (Nagel, 2010, p. 302)

If such a bias could affect our judgements involving epistemic modals, it might at least go some way towards explaining the data. The cases under discussion involve a more well-informed subject evaluating a claim made by a less informed subject.¹⁰ Perhaps Mary is evaluating John's claim as if he had the same information as she does. That would include the information that Harry is not in his office. If she did, that could explain why she responds in the way that she does. Moreover, when John is evaluating his past claim, it could be that he is subject to the same bias. He would then be evaluating his earlier assertion as if his earlier self had the same information as his current self. In that case, it would not be surprising if he took himself to have been wrong.

There is obviously a lot more that needs to be said and done before we can determine whether there is a real worry here. For one thing, there is the question of whether it is plausible that the bias is operative in the cases under discussion. There are also questions having to do with when we are affected by this kind of bias and whether there are circumstances in which we are more likely to overcome it. If there are situations in which the subjects are not affected by this kind of bias, what happens in those situations? Do we see different patterns of data those cases? In order to answer these question one would probably have to have a closer look at the relevant psychological research, and perhaps it is even necessary to carry out further research. That is beyond the scope of the present discussion, but it is something that deserves to be investigated further.

¹⁰For a discussion of so-called 'ignorant assessor' cases in which a less informed subject evaluates the assertion of a more well-informed subject, see e.g. Dietz (2008) and MacFarlane (forthcoming-a).

4.5 More on Disagreement and Retraction

Even if we acknowledge that the data needs to be taken seriously and cannot be dismissed as the result of a general psychological bias, it is not clear what conclusions we can draw about the semantics of epistemic modals. As bad as things might look at first glance, I do not think that we should be too quick to abandon contextualism. The issues involving disagreement and retraction are more subtle than they initially appear to be. I am going to start by looking at some potential worries about the disagreement data, and then I am going to argue that there are similar problems with the retraction data.

In order for the disagreement data to be a problem for contextualism, it needs to be shown that the disagreement targets the proposition expressed by the sentence that is being uttered. There are general reasons for thinking that this assumption is dubious. Disagreement is not always that simple.

Consider a situation in which Pierre and Marie are conducting a series of experiments in order to test a scientific hypothesis. They both believe the hypothesis to be false, but they have been unable to conclusively establish that it is false. After one of the experiments, the following conversation takes place:

- (55) a. Pierre: The hypothesis is false.
b. Marie: No, we need to do further testing.

There is some sort of disagreement going on in (55), but we do not want to say that Marie is denying that the hypothesis is false. That would be odd as long as we are assuming that she believes that the hypothesis is false. It is more plausible to take her to be disagreeing with Pierre on the grounds that there is insufficient evidence for him to assert that the hypothesis is false without conducting further experiments.¹¹

More importantly, there are more specific reasons to doubt that the disagreement in the cases under discussion must target the proposition expressed by the sentence uttered. As von Stechow and Gillies (2008) point out, in the case of first person uses of attitude verbs like ‘thinks’ or ‘believes’, the disagreement does not have to target the proposition expressed by the attitude report. To illustrate this point, we can start by considering the difference between (56bi) and (56bii).

¹¹ See e.g. von Stechow and Gillies (2008, p. 83) for a similar observation.

- (56) a. A: I think it's raining.
b. i. B: No, it isn't.
ii. B: No, you don't.

In the case of (56bi), B is denying that it is raining, not that A thinks that it is raining. The response in (56bi) differs from B's response in (56bii), which concerns the attitude report itself. It is convenient to talk about the disagreement as targeting a certain proposition. In the case of (56bi), the disagreement is targeting the proposition that it is raining. But in the case of (56bii), it is targeting the proposition that A thinks that it is raining. In the former case we can say that the disagreement is targeting the embedded proposition, and in the latter case we can say that it is targeting the attitude proposition.

Having made the observation that the disagreement can target the embedded proposition in cases like (56), von Stechow and Gillies go on to suggest that something similar may be going on in cases like (50).

- (50) a. John: Harry might be in the office.
b. Mary: No, he isn't. I just checked and he wasn't there.

Instead of interpreting Mary as denying the proposition expressed by (50a), we may interpret her as denying that Harry is in the office. If this is what Mary is doing, there is no problem for contextualism. In this case, we can talk about this in terms of the disagreement targeting the embedded proposition rather than the modal proposition. What the opponents of contextualism need to do, is to find a way of ensuring that the disagreement targets the modal proposition. However, that might prove difficult. In what follows, I will defend the hypothesis that we can understand the allegedly problematic cases of disagreement by taking the disagreement to be targeting the embedded proposition.

So far the discussion has concerned the disagreement data, but I want to argue that the retraction data is problematic in much the same way as the disagreement data. Again, the critical assumption is that the retraction targets the modal proposition. Since the corresponding assumption turned out to be dubious in the cases involving disagreement, it is perhaps not surprising that it turns out to be questionable when we are considering cases of retraction as well.

It is easy to come up with cases of retraction involving attitude verbs like 'thinks' that seem to be analogous to the cases of retraction involving 'might'.

- (57) a. John: I think that Harry is in the office.
b. Mary: No, he isn't. I just checked and he wasn't there.
c. John: Okay, then I guess I was wrong.

Again, it is quite natural for John to retract his original claim. His response in (57c) sounds perfectly appropriate in light of Mary's observation. Not only is his retraction appropriate, his attempt to stand by his original claim in (58c) is again very stilted and unnatural.

- (58) a. John: I think that Harry is in the office.
b. Mary: No, he isn't. I just checked and he wasn't there.
c. John: Okay, then he can't be in his office. But I still stand by what I said a second ago.

Having said that, I do not think that it is plausible to interpret John as retracting the attitude proposition. He presumably did think that Harry was in the office, and there is no apparent reason for him to retract the attitude report. Furthermore, this kind of example will probably not make us revise the semantics for 'thinks', say, along relativist lines. It is much simpler and much more plausible to interpret the retraction as targeting the embedded proposition. What John is retracting is the proposition that Harry is in the office, not the proposition that he thought that Harry was in his office.

The worry that is emerging is analogous to the problem facing attempts to argue against contextualism by appealing to cases of disagreement. Why should we be impressed by the retraction data involving epistemic modals when we get similar retraction data involving attitude verbs like 'thinks'? Why not conclude that the same thing is going on in both cases? Assuming that this is the correct way to think about cases like (57), that would amount to treating the retraction in the case of (51) as targeting the proposition that Harry is in the office.

- (51) a. John: Harry might be in the office.
b. Mary: No, he isn't. I just checked and he wasn't there.
c. John: Okay, then I guess I was wrong.

I am arguing that we should think of the disagreement and retraction data along similar lines. Just as it has been suggested that we should think of the disagreement as targeting the embedded proposition, I want to suggest that we should

think of the retraction as targeting the embedded proposition. In effect, I am putting forward the following hypothesis about what is going on in cases like (50) and (51):

The Hypothesis: In the cases under discussion, the disagreement and the retraction is targeting the embedded proposition.

In what follows, I will be discussing and defending this hypothesis. If the hypothesis is correct, it would allow contextualists to make sense of what is going on in the allegedly problematic cases of disagreement and retraction. Even if the modal proposition is true, as contextualists claim, the embedded proposition can still be an appropriate target of disagreement and retraction. It also provides the basis for an explanation of why it is also possible for the speaker to stand by her original claim. In that case, what is at issue is the modal proposition, not the embedded proposition. What she is defending is the truth of the modal proposition that she asserted. If contextualism is correct, that proposition is true. It is therefore natural for her to be in a position to defend her assertion on that basis.

One way of illustrating the point is by considering the following argument:

1. If a contextualist semantics for epistemic modals is correct, cases like (50) and (51) are not cases of appropriate disagreement and retraction.
 2. Cases like (50) and (51) are cases of appropriate disagreement and retraction.
- ∴ 3. A contextualist semantics for epistemic modals is not correct.

Perhaps this argument is a bit crude, and I am not claiming that this is what opponents of contextualism actually have in mind, nor that this is the best way of stating the case against contextualism. Be that as it may, the argument is still useful for the purpose of illustrating the point I am making. The second premise can be seen as motivated by taking our judgements about cases like (50) and (51) at face value. This is not beyond dispute, but I want to focus on the motivation for the first premise. One reason for thinking that a contextualist semantics is incompatible with there being disagreement and retraction in the cases under discussion is that the modal proposition expressed by the first sentence uttered by

John is such that it would be strange to think that the disagreement and retraction could be targeting that proposition. Thinking along these lines, one could offer the following argument for the first premise of the original argument:

- 1.* If a contextualist semantics for epistemic modals is correct, there is no disagreement and retraction in (50) and (51) that target the modal proposition.
 - 2.* If there is no disagreement and retraction in (50) and (51) that target the modal proposition, cases like (50) and (51) are not cases of appropriate disagreement and retraction.
- ∴ 3.* If a contextualist semantics for epistemic modals is correct, cases like (50) and (51) are not cases of appropriate disagreement and retraction.

Again, neither the first nor the second premise of this argument is beyond dispute, but it is the second premise of this argument that I will focus on. I want to argue that even if we grant the first premise of this argument, we have good reasons to reject the second premise. If we allow for the possibility of disagreement and retraction that target the embedded proposition and not the modal proposition, it is possible for there to be disagreement and retraction in (50) and (51) even if contextualism is true and the disagreement and retraction cannot target the modal proposition. Furthermore, once we allow for the possibility of disagreement and retraction that do not target the modal proposition, we should be hesitant to accept the first premise of the original argument. Even if a contextualist semantics is correct, and even if that rules out that there is disagreement and retraction that target the modal proposition, there can still be disagreement and retraction targeting the embedded proposition. We can thus view the hypothesis that the disagreement and retraction are targeting the embedded proposition as offering a way of resisting this argument by undermining the motivation for the first premise of the original argument.

What the case of ‘thinks’ shows, is that it is sometimes possible to understand disagreement and retraction as targeting an embedded proposition. I have been assuming that there is a relevant sense in which the behaviour of attitude verbs can cast light on the behaviour of epistemic modals. If it can be shown that there are reasons for thinking that attitude verbs like ‘thinks’ and epistemic modals

should be treated differently with respect to disagreement and retraction, that would cast doubts on the hypothesis. For one thing, it is not obvious that attitude verbs like ‘thinks’ provide the best point of comparison for epistemic modals like ‘might’. In what follows, I will discuss further issues and try to answer some potential objections along these lines. In particular, I will discuss some relevant points that have been raised by Stephenson (2007) and MacFarlane (forthcoming-a). In the end, I will argue that the issues I consider do not show that we need to treat attitude verbs and epistemic modals differently with respect to the relevant issues involving disagreement and retraction. I will also argue that we find similar patterns of data in a fairly wide range of cases, and not just in the case of epistemic modals like ‘might’ and attitude verbs like ‘thinks’.

4.6 Stephenson on Disagreement

In the previous section, I followed von Fintel and Gillies in pointing out the similarities between attitude verbs like ‘thinks’ and epistemic modals when it comes to disagreement data. But this is not the only useful point of comparison for epistemic modals. Indeed, there are other constructions in English that one would expect to be even more relevant in this respect. Stephenson (2007) raises an issue concerning contextualist treatments of epistemic modals along such lines. She argues that if contextualism were true we might expect Mary’s response in (59b) to be just as appropriate as her response in (50b).

- (50) a. John: Harry might be in the office.
 b. Mary: No, he isn’t. I just checked and he wasn’t there.
- (59) a. John: I don’t know that Harry is not in his office.
 b. Mary: No, he isn’t. I just checked and he wasn’t there.

But there is a clear difference between the two dialogues. Whereas Mary’s response in (50b) is perfectly natural, there is something awkward about Mary’s response in (59b). Stephenson notes that this is something that requires explanation. If we assume that the simple contextualist treatment of epistemic modals is correct, (59a) and (50a) will have the same truth-conditions since the latter is true just in case it is not ruled out by what John knows that Harry is in the office. The upshot of this is that someone who defends a contextualist semantics for

epistemic modals needs to explain both the presence of disagreement in (50) and the contrast between that case and (59).

However, it is not clear what the difference between (50) and (59) really shows. If we change the example slightly, we seem to get a different verdict.

- (60) a. John: For all I know Harry is in the office.
b. Mary: No, he isn't. I just checked and he wasn't there.

Again we find a significant difference between the cases. Whereas (59b) is admittedly awkward, (60b) is much better. But then it is just as much a question of how we can explain the difference between (59) and (60) as it is a question of how we can explain the difference between (50) and (59). When looking at the question in this way, it is far from obvious why this is only a question for contextualists.

However, I also want to gesture towards an explanation of this difference that fits well with the hypothesis that the disagreement can be explained by taking the disagreement to be targeting the embedded proposition. It is worth observing certain differences between (50) and (60) on the one hand, and (59) on the other. The complement clause in (59a) is different from the complement clauses in (50a) and (60a) insofar as the former is negated, whereas the latter are not. This difference is made explicit in (61) and (62).

- (61) It might be that [_{CP}Harry is in the office].
(62) I do not know that [_{CP}Harry is not in the office].

If we assume that what is going on in these cases is that the disagreement is targeting the embedded proposition, this makes it less surprising that the cases are different. In order for the disagreement to be targeting the embedded proposition in the case of (59), that is, the proposition expressed by the complement clause, Mary would have to deny that Harry is not in his office, but that is not what she is doing. Insofar as she is denying anything, she is denying that Harry is in the office. It is therefore not surprising that her response in (59) sounds strange if we interpret her as trying to deny that Harry is not in his office.

At worst these considerations appear to be inconclusive with respect to the prospects of holding on to a contextualist semantics for epistemic modals, but it is also tempting to draw a more optimistic conclusion. For one thing, it looks

like the data pattern in a way that one would expect if the hypothesis that the disagreement is targeting the embedded proposition is correct. It is also interesting to observe the difference between (59) and (60). Cases like (60) will turn out to be important in the discussion that follows.

4.7 The Disagreement Constraint

MacFarlane (forthcoming-a) also points to disagreement data as a motivation for abandoning contextualism in favour of relativism. He argues that a conversation could be centred around the question of whether Joe might be in Boston rather than the question of whether Joe is in Boston. As MacFarlane puts it:

A conversation might center, for a time, on the question whether Joe might be in Boston. The issue is not whether Joe *is* in Boston; everyone present acknowledges that he might be in Berkeley, and so no one thinks that there are going to be grounds for asserting that he *is* in Boston. The point of the conversation is to settle whether he *might be* in Boston. Reasons are offered on both sides, disputes are resolved, and perhaps a consensus is reached. (MacFarlane, forthcoming-a, p. 6)

In this conversation, nobody would think that they are going to establish that Joe is in Boston since they cannot rule out that he is in Berkeley. But we could still find apparent disagreements between the participants as they debate Joe's whereabouts. Suppose that Mary and John are participants in the conversation. Whereas Mary knows that Joe is in California, it is compatible with what John knows that Joe is in Boston. In that case, it would be natural for an exchange like the one in (63) to take place.

- (63) a. John: Joe might be in Boston.
b. Mary: No, he can't be. He isn't allowed to leave the state.

However, even if consider a conversation in which the participants recognise that they cannot rule out that Joe is in Berkeley, we should not be too quick to assume that the disagreement concerns the modal proposition. It is not clear what MacFarlane has in mind when he says that the conversation is centred around the

question of whether Joe might be in Boston. Presumably the conversational participants are still mainly interested in Joe's whereabouts, and they may be able to rule out certain possibilities without being able to rule out certain other possibilities. Even if nobody thinks that they are going to establish that Joe is in Boston, they can still think that it is possible to establish that Joe is not in Boston.

We can think of what is going on in such a conversation as an attempt to narrow down a set of locations.¹² It is therefore relevant whether a particular participant in the conversation can or cannot rule out that Joe is in a certain location, and it would also be relevant if someone could offer reasons for thinking that Joe is or is not in a certain location. We should therefore be careful about reading too much into MacFarlane's description of the conversation as one that is centred around the question of whether Joe 'might' be in Boston. For all that MacFarlane has said, we can still take the conversational participants to be primarily interested in Joe's whereabouts and the proposition that Joe is in Boston.

However, there is an objection that can be raised at this point. An important feature of the conversation that MacFarlane describes, is that it is reasonable to assume that the conversational participants do not take themselves to be in a good enough position to believe or assert that Joe is in Boston. After all, they take themselves to be unable to rule out the possibility that Joe is in Berkeley. In particular, Mary cannot assume that John believes that Joe is in Boston when he utters (63a). This is relevant if we adopt the following necessary condition on disagreement:

The Disagreement Constraint: Disagreement can only target a proposition that is believed or asserted.

I do not want to attribute this principle to MacFarlane, but it is nevertheless interesting to discuss. It makes it problematic to argue that what is going on in the conversation under discussion is that the disagreement in (63) concerns the proposition that Joe is in Boston. If we are assuming that John does not believe or assert that Joe is Boston, we cannot say that the disagreement targets that proposition. In fact, if the disagreement constraint is correct, it is hard to see how one

¹²This is more or less how we would think about the case if we adopt Stalnaker's (1978) picture of assertion.

could hope to explain all the allegedly problematic disagreement data by arguing that the disagreement targets the embedded proposition. Presumably, we are often prepared to make a modal claim like (63a) without wanting to either assert or believe the embedded proposition.

The question is whether the disagreement constraint is plausible. Consider a case in which the speaker is just making a guess and presumably neither asserts or believes the proposition in question. For instance, suppose that A, B, and C are watching a card game. While they are waiting for the dealer to reveal the next card, A and B are trying to guess whether the next card will be a spade, a heart, a diamond or a club.

- (64) A: Can you guess which suit the next card will be?
a. B: I don't know. I am guessing that next card will be a spade.
b. C: No, it won't. The deck is stacked and the next card is the queen of hearts.

If we assume that C is right about what the next card is, his response sounds fine. What he is denying is the proposition that the next card will be a spade. No other obvious candidate is available. He is surely not denying that B is guessing that the next card will be a spade, regardless of whether that proposition is even asserted or not. However, since B was just guessing, we cannot assume that he believes that the next card will be a spade. Understood in those terms, (64) appears to be a counterexample to the disagreement constraint. Moreover, if a mere guess will do, it is not at a surprise that we get similar data involving epistemic modals. Of course, someone could contest that this should be described as a case of disagreement, but then it is not clear that we need to describe the cases involving epistemic modals as cases of disagreement either. Describing the case one way or another does not solve the problem.¹³

To make matters even more difficult for the opponents of contextualism, it was observed in the previous section that we can have an exchange like the one

¹³I do not take myself to be committed to treating guessing as a *sui generis* speech act, but if one wants to say that cases like (64) are actually cases of assertion in order to get around this objection to the disagreement constraint, I would also be more inclined to say the same thing about a case in which someone is using an epistemic modal.

in (6o) even if one of the speakers makes it explicit that she is talking about what is compatible with what she knows.

- (6o) a. John: For all I know Harry is in the office.
b. Mary: No, he isn't. I just checked and he wasn't there.

If the disagreement constraint is correct, we cannot understand the disagreement as targeting the embedded proposition, since John does not believe or assert that proposition. But (6oa) is not very amenable to a relativist treatment. The point about (6oa) is precisely that John makes it explicit that he is talking about what is compatible with what he knows. From that point of view, we seem to be much better off treating Mary as denying that Harry is in the office. However, then it is strange to think that this could not happen in (63) as well, and to claim that we need to view cases like (63) as providing evidence for some kind of alternative semantics for epistemic modals.

The point here is mainly negative. We should reject the disagreement constraint in favour of some less stringent requirement on disagreement. A further question is what the replacement for the disagreement constraint is going to be. What examples like (64) and (6o) indicate, if anything, is that the requirement will not be very strict, and there is no reason to think that a plausible requirement will rule out the possibility of disagreement targeting the embedded proposition in cases involving epistemic modals.

4.8 The Retraction Constraint

So far I have been focusing on the issues having to do with the disagreement data. However, similar issues arise in the case of retraction data. This sort of data is also central to MacFarlane's case against contextualist treatments of epistemic modals. In particular, he offers specific reasons for thinking that cases like (51) and (57) are different, and that the retraction in cases like (51) cannot target the embedded proposition.

- (51) a. John: Harry might be in the office.
b. Mary: No, he isn't. I just checked and he wasn't there.
c. John: Okay, then I guess I was wrong.

- (57) a. John: I think that Harry is in the office.
b. Mary: No, he isn't. I just checked and he wasn't there.
c. John: Okay, then I guess I was wrong.

He points out that in the case of (51), John does not assert that Joe is in the office, and we can easily imagine a situation in which he does not believe it either. In this sense, (51) is different from (57). In the latter case, it is at least safe to assume that he believes that Joe is in the office.¹⁴ MacFarlane seems to think that this shows that he cannot be retracting the proposition that Joe is in the office in (51). It would seem that in order for this to rule out that the retraction targets the embedded proposition, something like the following necessary condition on retraction needs to be correct:

The Retraction Constraint: Retraction can only target a proposition that is believed or asserted.

Of course, as MacFarlane is aware, we sometimes retract other speech acts than assertion, and it is therefore natural to think that this is not a very plausible necessary condition for retraction.¹⁵ However, it is not always clear that the retraction involves a proposition in those cases, and it is also important that we are mainly interested in a fairly specific form of retraction data following the pattern in (51) and (57).

But even setting aside these general concerns, how plausible is the retraction constraint? The retraction constraint is the analogue of the disagreement constraint for retraction, and as we saw in the previous section, the disagreement constraint turned out to be dubious. I will argue that there are also good reasons to think that the retraction constraint does not hold as a necessary condition on retraction. This will hopefully be clear once we look at some specific cases. Consequently, one cannot simply claim that John cannot be retracting the proposition that Joe is in the office because he does not believe or assert that proposition.

¹⁴There is a further question about whether John asserts that Joe is in the office by uttering (57). I suspect that the answer will depend on how one thinks about assertion, but I am not going to pursue this question further at this point.

¹⁵MacFarlane (2011) points out that we sometimes retract questions, commands, and apologies.

Pretend that Holmes and Watson are investigating some heinous crime. Watson harbours certain suspicions towards Smith, but he does not yet believe that Smith is the culprit, as there are many other suspects. He has still not made up his mind. As a matter of fact, there is also certain strong evidence in favour of Smith's innocence that he has not yet uncovered. Watson proceeds to discuss the matter with Holmes who has had the opportunity to review the evidence.

- (65) a. Watson: I have a suspicion that Smith did it.
b. Holmes: No, he couldn't have done it. He was in the Far East.
c. Watson: Okay, then I guess I was wrong.

It seems appropriate for Watson to retract his original claim in light of the evidence presented by Holmes. I am assuming that in doing so he is retracting the proposition that Smith did it, and not the proposition that he had a suspicion that Smith did it. After all, that proposition is true, and that was not what Holmes was challenging. However, we were also taking it to be a part of the scenario that Watson did not believe that Smith did it. He only had a suspicion towards Smith, and I am assuming that suspicion is compatible with lack of outright belief. The retraction constraint incorrectly predicts that Watson cannot be retracting the proposition that Smith did it.

Again, it also makes sense to look at cases in which the speaker is merely making a guess. Consider the following dialogue:

- (66) a. A: Can you guess when Napoleon was born?
b. B: I don't know. I am guessing that he was born in 1768.
c. A: No, he wasn't. He was born in 1769.
d. B: Okay, then I was wrong.

The response in (66d) sounds appropriate, but since B is only making a guess we cannot assume that B either believes or asserts that Napoleon was born in 1768 on the basis of her utterance of (66a). Moreover, it would clearly be odd for B to stand by her guess once she is appropriately corrected. If the retraction constraint is correct, we cannot say that what A is retracting is the proposition that Napoleon was born in 1768. But then how do we make sense of her retraction?

Perhaps we are better off if we give up the original formulation of the retraction constraint in terms of belief and assertion, and instead try to state a weaker

necessary condition on retraction that allows for cases in which the speaker is performing a speech act other than assertion or has an attitude other than belief towards the relevant proposition. This may well be the right thing to do, but it is not clear that this will do the work that the original constraint was intended to do in terms of ruling out the possibility that John is retracting the embedded proposition in (51). The worry is that a weaker constraint does not allow us to distinguish the retraction cases under discussion from the retraction cases involving epistemic modals. This sort of worry is even more pressing if one takes John's retraction in (67) to be appropriate.

- (67) a. John: For all I know Harry is in the office.
b. Mary: No, he can't be. I just checked and he wasn't there.
c. John: Okay, then I guess I was wrong.

As far as I can see, it does not sound altogether unnatural for John to retract his original claim. But there is no more reason to think that John believes or asserts that Joe is in the office on the basis of his utterance of (67a) than there is in the case of (51).

All of this seems to amount to a fairly strong case against the retraction constraint. It is not necessary for a proposition to be the target of retraction that the proposition is asserted or believed. In light of these considerations, it looks like it is not enough to establish that the speaker does not believe or assert the embedded proposition in order to show that the retraction in cases like (51) cannot be targeting that proposition.

4.9 Disagreement and Credences

So far the point of the discussion has been mainly negative. The requirement that a proposition must be believed or asserted in order for it to be the target of disagreement or retraction is too strong. However, we are still left with the question of why the disagreement and retraction can target the embedded proposition in the cases under discussion. What makes the embedded proposition available as a target for disagreement and retraction? If it is not the fact that the proposition is believed or asserted, then it must be something else. In this section the focus will be on disagreement. Insofar as the disagreement cannot be understood in

terms of what the parties believe, a natural thought is that it has something to do with the degrees of belief or credences that the parties have in the embedded proposition. The idea is that if two parties have different credences in a proposition, there is a sense in which their doxastic states are in conflict and they can be said to disagree. This is captured by (68).

(68) Two parties disagree if there is a proposition such that they have different credences or degrees of belief in that proposition.

As a sufficient condition for disagreement, (68) is likely to be somewhat controversial.¹⁶ To illustrate the idea, consider a situation in which two detectives are investigating a murder. The only suspect is Smith, but the detectives have not made up their minds as to whether he is the murderer or not. They have not formed the outright belief that he is the murderer, but they have not formed the outright belief that he is not the murderer either. However, the first detective has a higher credence in the proposition that Smith is the murderer than the second detective. If we represent credences numerically, we can represent the first detective as having a credence of .7 in the proposition that Smith is the murderer, while the second detective has a credence of .6 in the proposition that Smith is the murderer.¹⁷ This difference will most likely be reflected in their behaviour in various ways. For instance, the detectives might differ with respect to the bets they are willing to make, and with respect to how they react to new evidence. In this case, it is natural to regard the difference in credences as sufficient for the detectives to disagree.¹⁸

¹⁶See e.g. Cappelen and Hawthorne (forthcoming) who seem reluctant to talk about disagreement when it comes to degrees of belief or credences. They note that it is more natural to talk about disagreement as applying to beliefs rather than weighted credences, but beyond that they do not offer any further arguments. However, in ordinary speech it is rare to find speakers talking explicitly about their credences, so it is not clear that this is a major concern.

¹⁷I am assuming that having a relatively high credence, such as .7, is compatible with not having an outright belief in the relevant proposition. For the purpose of the following discussion I am also assuming that credences are precise, though as far as I can see, nothing here turns on whether we allow credences to be imprecise. See e.g. Joyce (2010) and White (2010) for recent discussion of imprecise credences.

¹⁸In a realistic scenario, it will typically be the case that if two parties have different credences in a proposition, they also have different beliefs regarding related propositions. This makes it harder to get clean judgements about whether the disagreement is a result of a difference in credences.

An immediate worry with this view is that it predicts too much disagreement. If the slightest difference in our credences entail that we disagree, that might be seen as problematic. While I recognise that someone might be worried about this, it is not clear that this is an unacceptable consequence. It is true that insofar as the differences in credences are too small to be noticeable, we are not likely to think of the parties as disagreeing. However, that is something we might be able to live with. After all, there might be all sorts of propositions that we disagree about without anyone ever noticing. That is true regardless of whether we allow a difference in credences to constitute a disagreement. Furthermore, when the differences are noticeable, it seems natural to think that there is a sense in which the parties disagree.

If one is still worried about predicting too much disagreement, a possibility is to say that we have disagreement only if there is a significant difference in credences. What counts as 'significant' will presumably depend on the context. While I am not in principle opposed to introducing an additional element of context-dependence or vagueness, a worry is that it makes it harder to determine whether the relevant condition is satisfied. In the interest of keeping things as simple as possible, I will therefore assume that (68) is adequate as it stands. For the purpose of the present discussion, this ought to be sufficient, but a more realistic view might be one that introduces an additional element of context-dependence or vagueness.

If one wants to use (68) to explain cases of disagreement involving epistemic modals, we also have to assume that the parties have different credences in the relevant propositions. In the case of (48) and (49), this seems to be a plausible assumption.

(48) The person approaching might be Fred.

(49) The person approaching cannot be Fred.

In this case one would probably expect the speaker of (48) and the speaker of (49) to have different credences in the proposition that the person approaching is Fred. However, if one is looking for additional reasons for thinking that the speakers have different credences in the relevant proposition, one possibility is to exploit the connection between beliefs, credences, and knowledge. The speaker of (48) takes it to be compatible with what she knows that the person approaching

is Fred, while the speaker of (49) takes this to be ruled out by what she knows. If one thinks that speakers ought to adjust their credences based on what they know, and one assumes that the speakers are trying to observe the relevant rule, this would give us some grounds for thinking that their credences will be different.¹⁹ This line of argument requires some fairly controversial assumptions about the relationship between knowledge, belief, and credences, but the point is mainly to give a rough idea of how one can provide evidence for a difference in credences even though that is not entailed by a contextualist semantics for epistemic modals.

We can also understand the disagreement between Mary and John in (50) as based on a difference in credences.

- (50) a. John: Harry might be in the office.
b. Mary: No, he isn't. I just checked and he wasn't there.

We can account for the disagreement between Mary and John in terms of Mary and John having different credences in the proposition that Harry is in the office. It has already been observed that it is implausible to hold that it is necessary for the disagreement to be targeting the embedded proposition that John believes that Harry is in the office. Instead, the proposal is that it is sufficient that John has a certain positive credence in the embedded proposition and that he conveys that by his utterance of (50a). This is a fairly weak condition, but in light of the preceding discussion this seems appropriate.

It is useful to compare contextualism to a view that Yalcin (2007, 2010, forthcoming) calls 'credal expressivism'. A simple way of thinking about this view is to say that when a speaker utters a sentence containing an epistemic modal, she expresses her credence in the embedded proposition. In the case of 'might' this could simply be a positive credence in the embedded proposition, whereas in the case of 'must' it would probably be a high credence in the embedded proposition. In the case of (50), John would express his positive credence in the proposition that Harry is in the office, while Mary would express a high credence in the proposition that Joe is not in his office. A credal expressivist could then argue

¹⁹See (Williamson, 2000) for relevant discussion. If one accepts the view defended by Williamson that one's evidence consists in what one knows, it is natural to think that people ought to try to adjust their credences based on what they know.

that the reason why we take Mary and John to disagree is that their credences in the relevant propositions are different.

The contextualist does not treat epistemic modals as having such a direct connection to credences. It is not a part of the contextualist picture that speakers express credences when using epistemic modals. However, what is relevant for our purposes is that contextualists can tell more or less the same story about disagreement as credal expressivists.²⁰ It does not matter whether the speakers express the relevant attitudes in the way that credal expressivists are talking about. What matters is what credences the speakers actually have. In the case of (50), it seems plausible that Mary and John have different credences in the proposition that Harry is in the office. For our purposes, that should be enough.

The upshot of this is that if something like (68) is correct, it offers a basis for thinking that the speaker of (48) and the speaker of (49) disagree without committing one to the idea that this must be cashed out in terms of a conflict of outright beliefs. The basic idea is to understand the disagreement in terms of a difference in credences and not in terms of a difference in beliefs.

4.10 The Pragmatics of Epistemic Modals

In the previous section I dealt with issues involving disagreement. In this section I will focus on issues involving retraction. While it is doubtful whether a speaker must assert or believe a proposition in order for retraction to be appropriate, a possible rejoinder is that there must be some story about why it is appropriate for a speaker to retract a proposition. So far I have been talking about retraction as targeting a proposition, but it is plausible to think that of retraction as targeting a speech act.²¹ A natural question is whether sentences containing epistemic modals can be used to perform a speech act involving the embedded proposition, in addition to being used to assert the modal proposition, and whether this can shed some additional light on the retraction data. Insofar as we want to connect what is going on at the level of mental states with what is going on at the level of

²⁰Dreier (1999) makes a similar point with respect to expressivism in moral philosophy. Insofar as an expressivist is in a position to tell a plausible story about disagreement, it is not clear why the contextualist cannot tell more or less the same story. See e.g. also Jackson and Pettit (1998).

²¹See e.g. MacFarlane (2011).

language and discourse, this can also contribute towards a better understanding of the disagreement data.

The idea is that there are two speech acts performed, one involving the assertion of the modal proposition and another speech act involving the embedded proposition. Proposals along these lines have been put forward by von Stechow (2007, 2008) and Portner (2009). My goal here is not to settle on a definitive view, but rather to sketch what I take to be live options and indicate promising avenues of further investigation.

The view that epistemic modals can be used to perform two distinct speech acts has certain *prima facie* promising features. Insofar as it is sometimes possible to appropriately stand by one's original assertion, this is not surprising if there are two distinct speech acts involved. In those cases, it is the assertion of the modal proposition that is at issue. However, we often focus on a speech act involving the embedded proposition. In the terminology of Simons (2007), it is the 'main point' of the utterance. That is what the retraction is targeting when the speaker chooses to retract.

This kind of view might also give us a better grip on examples like (54), in which retraction seems inappropriate.

- (54) a. John: Joe might be coming to the party, but it would be very surprising if he did.
b. Mary: Actually he won't be coming to the party. I just talked to him a minute ago.
c. John: Okay, then I guess I was wrong.

It is not entirely clear how relativists can deal with examples like this. Perhaps we can do better if we adopt a view on which sentences containing epistemic modals can be used to perform two distinct speech acts. A natural suggestion is that adding 'but it would be surprising if he did' results in the relevant speech act involving the embedded proposition being canceled. In that case, it would not be possible for the retraction to be targeting that speech act. But if contextualism is correct, we already have reasons for thinking that it would not be appropriate to retract the assertion of the modal proposition. These points suggest that there is something to be said for exploring whether sentences containing 'might' can be used to perform a speech act involving the embedded proposition in addition

to an assertion of the modal proposition.

More generally, it is important not to lose sight of the fact that there are many situations in which it is not appropriate for a speaker to retract a speech act, even if the proposition in question turns out to be false. The cases discussed in the previous section indicate that there is a fairly wide range of cases in which it sounds natural and appropriate for a speaker to retract in light of new information, even if the speaker does not believe or assert the proposition in question. But it is also important to look at situations in which it is not natural and appropriate for a speaker to retract. This indicates that, say, raising a possibility to salience cannot be sufficient to make a proposition an appropriate target for retraction. For instance, MacFarlane (forthcoming-a) offers the following example:

- (6g) a. A: It is rumored that you are leaving California.
b. B: That's completely false!
c. A: Okay, then I was wrong. I take back what I said.

We can imagine that A has heard a rumour to the effect that B is leaving California, but that he has not made up his mind as to whether he should believe the rumour or not. In that case, A's retraction sounds odd. Assuming that the possibility that A is leaving California was not already salient, it will surely be made salient by A's utterance of (6ga). Since it would be strange for A to retract once it turns out that A is not leaving California, the mere fact that a speaker is raising the possibility that his interlocutor is leaving California to salience, cannot be sufficient to make it appropriate for the speaker to retract the previous speech act if it turns out that this is false. Insofar as it is not sufficient in the case of (6g), we cannot assume that it is sufficient in the case of (51) either.

- (51) a. John: Harry might be in the office.
b. Mary: No, he isn't. I just checked and he wasn't there.
c. John: Okay, then I guess I was wrong.

This is something that we need to take into account when we evaluate different proposals regarding what kinds of speech acts epistemic modals can be used to perform. I will look at suggestions put forward by von Stechow and Gillies (2007), Portner (2009), and Swanson (2006). The main goal is to achieve a better understanding of our options, and to identify the options that look promising from

the point of view of making sense of the retraction data.

4.10.1 Proffering

One proposal that has been put forward by von Stechow and Gillies (2007), following a similar suggestion made by Simons (2007), is that epistemic modals can be used to perform two speech acts. The first is an assertion of the modal proposition as on the standard possible worlds semantics for epistemic modals. The second is a proffering or putting forward of the embedded proposition. Admittedly, von Stechow and Gillies do not say much about what proffering amounts to in this context. While the basic idea is intuitive, it is not clear whether their proposal provides the resources to explain the appropriateness of retraction in cases involving epistemic modals. To illustrate the need for clarification, consider the following example from Simons (2007):

- (70) a. A: Who was Louise with last night?
b. B: Henry believes that she was with Bill.

Simons proposes that B can be understood as proffering that Louise was with Bill last night as an answer to A's question in (70). This is a natural thing to say, but consider how strange the following continuation of the dialogue would be:

- (71) a. A: No, Louise couldn't have been with Bill. He was in Princeton.
b. B: Okay, then I guess I was wrong.

There is surely something odd about B's retraction in (71b). If Simons is right to describe B as proffering that Louise was with Bill as an answer to A's question, it does not suffice to say that John was proffering that Harry is in the office in order to explain why it is appropriate for him to retract and say that he was wrong once it turns out that Harry was not in the office.

This result should not be seen as too discouraging though. Proffering might involve varying degrees of commitment to the truth of the proposition in question, and the commitment might be weaker when the speaker is explicitly reporting on a rumour or what someone else thinks. Simons thinks that (72a)-(72f) are ways of proffering that Louise was with Bill as an answer to the question of who Louise was with last night, but in that case the speaker is surely conveying varying

degrees of certainty and commitment with respect to the proposition that Louise was with Bill.

- (72) Who was Louise with last night?
- a. I'm convinced that she was with Bill.
 - b. I think that she was with Bill.
 - c. I imagine that she was with Bill.
 - d. I heard that she was with Bill.
 - e. Henry said that she was with Bill.
 - f. Henry suggested that she was with Bill.

It would be nice to have a story about the varying degrees of commitment involved in the different ways of proffering, but that is of more general interest and not a specific problem involving epistemic modals. I am not going to offer such a story here, but it is at least plausible that the cases in which the speaker is making it explicit that she is relying on a rumour or what someone else says or thinks are cases in which the speaker is conveying a weaker commitment to the embedded proposition. Simons argues that in cases like (72e), the speaker is conveying that it is Henry's belief that is the source of the proffered answer. Therefore it is reasonable to conclude that the speaker is not taking responsibility for the answer in the same way as if she was using (72a)-(72c).

4.10.2 The Common Propositional Space

Portner (2009) attempts to offer a more worked out proposal along more or less the same lines as von Stechow and Gillies. According to the familiar picture of assertion developed by Stalnaker (1978), the effect of an assertion is to add a proposition to the common ground. For our purposes, we can think of as the set of propositions being presupposed by the conversational participants.²² In addition to this familiar notion of a 'common ground', Portner introduces the notion of a 'common propositional space'. He characterises the common propositional space as 'the set of propositions in which the participants in the

²²Note that I will continue to talk in terms of the common ground, the set of propositions that are being presupposed, rather than the context set, the set of worlds compatible with what is being presupposed.

conversation are mutually interested' (2009, p. 175). On the picture he puts forward, the common ground is a subset of the common propositional space. The common propositional space also includes propositions that are candidates for inclusion in the common ground.

Portner suggests that we can use the notion of a common propositional space to understand the behaviour of epistemic modals. Like von Fintel and Gillies, Portner argues that an epistemic use of a sentence involving 'might' can be used to perform two distinct speech acts. It can be used to add the modal proposition to the common ground, and to add the modal proposition and the embedded proposition to the common propositional space. In other words, the modal proposition is being asserted as one would expect, but the embedded proposition is also being added to the common propositional space.

While Portner's proposal looks interesting, there are still certain aspects of the view that are not entirely clear, and this makes it hard to evaluate. In particular, it is not clear to what extent adding a proposition to the common propositional space involves more than merely raising it to salience. Assuming that it is not sufficient to license retraction that the speaker was raising a possibility to salience, this means that it is not clear whether Portner's proposal puts us in a position to explain the retraction data. This is not a criticism of the view. It is merely to point out that more might be needed if we want to make use of this view in order to explain the retraction data.

4.10.3 Doxastic Advice

What I take to be the most interesting and promising possibility is to follow Swanson (2006), and to think of the speech act involving the embedded proposition as a kind of doxastic advice.²³ This possibility is also entertained by von Fintel and Gillies (2007). Simplifying somewhat, we can say that an epistemic use of a sentence containing 'might' conveys the advice not to overlook the possibility

²³It is important to note that the overall treatment of epistemic modals defended by Swanson (2006) is different from the one that I am defending. However, I am not going to delve into the details of his probabilistic semantics. Even though Swanson is engaged in a more ambitious project of offering an alternative to a traditional semantics for epistemic in terms of quantification over possible worlds, it is not clear why someone working with a more standard possible worlds semantics could not take certain aspects of his view on board.

that the embedded proposition is true, and to have a certain positive credence in the proposition in question. I take it that we have at least some independent grip on the notion of giving advice, though for theoretical purposes one might hope that it will ultimately be possible to offer a more informative and rigorous characterisation.

Can this give us any traction with respect to the retraction data? There is some reason for optimism here, since we are routinely prepared to retract advice if we gain new and relevant information. Let us suppose that you are looking for advice about where to go on holiday and I advise you to go to somewhere based on certain information I have. I later acquire new information that makes it clear that the advice I gave you was bad, and that it would be a bad idea to go where I advised you to go. In that case, it would be appropriate for me to retract my advice, and to say that I was wrong. As Swanson points out, we do not have to say that I am subject to criticism in the sense that I could have been doing my best given the information that was available to me. Still, it seems appropriate for me to retract my previous advice in light of the new information that has been made available to me.

Another nice feature of this account is that it connects well with the account of disagreement that was discussed in §4.9. If we think of the disagreement in the relevant cases in terms of the parties having different credences, it would make sense to think that this also plays a role when it comes to retraction. The current proposal does a good job in this respect. What is going on in the case of (51), is that John is advising Mary to have a certain positive credence in the proposition that Harry is in the office. Once Mary provides him with new information regarding Harry's whereabouts, he retracts the advice.

This is not yet to show that this is what is going on in the case of epistemic modals, but I think it looks like a promising area for further research.²⁴ In any case, this is not meant to be the final word on the matter. My goal has been

²⁴A worry is that the claim that the speaker is advising her interlocutor to have a certain minimal credence in the embedded proposition is too strong. The proposal would make certain uses of 'might' inappropriate. However, in many cases in which the speaker knows that her interlocutor is better informed than she is, it is not appropriate to use an epistemic modal like 'might'. For instance, in the case of (6g), it sounds awkward if we replace 'It is rumored that you are leaving California' with 'You might be leaving California'. This is something we can explain if epistemic modals are used to give doxastic advice.

to outline various ways of developing the idea that epistemic uses of sentences involving ‘might’ can be used to perform more than one speech act. While it is to some extent an open question whether this helps us to explain the retraction data, both an account in terms of proffering as suggested by von Fintel and Gillies (2007) and an account in terms of giving advice inspired by Swanson (2006), look promising. Perhaps one way to think about the retraction data, is that it does not really teach us anything significant about the semantics of epistemic modals. Instead it teaches us something about the pragmatics of epistemic modals in the sense that it teaches us something about the speech acts that we perform when we use epistemic modals.

4.11 Further Problems

So far I have focused on considerations involving disagreement and retraction. The basic response I have offered on behalf of contextualism is to say that the disagreement and retraction target the embedded proposition. However, there are reasons to think that considerations involving agreement cannot be straightforwardly handled in the same way. The worry can be illustrated by an example discussed by von Fintel and Gillies (forthcoming). Let us suppose that Alex and Billy are looking for Billy’s keys when the conversation in (73) takes place.

- (73) a. Alex: You might have left them in the car.
b. Billy: You’re right. Let me check.

If the modal proposition is a proposition about what is compatible with Alex knows, it is not clear that it makes sense for the agreement to be targeting that proposition. According to von Fintel and Gillies, Billy is not in a position to comment on what is compatible with what Alex knows. However, it is also not clear that it makes sense for the agreement to be targeting the embedded proposition. It does not look like Billy agrees that the keys are in the car. After all, we may stipulate that she does not believe that the keys are in the car. This makes examples like (73) different from an examples involving disagreement like (50).

- (50) a. John: Harry might be in the office.
b. Mary: No, he isn’t. I just checked and he wasn’t there.

In the case of (50) it makes sense to interpret Mary as denying that Harry is in the office. It would be less plausible to stipulate that she does not believe that Harry is not in the office.

It is important to note that the current objection does not assume the correctness of the disagreement constraint or an analogous constraint on agreement. Instead, it rests on the assumption that a speaker must believe a proposition in order to agree with it. A similar constraint on disagreement would be the constraint that a speaker must believe the negation of a proposition in order to disagree with it. Insofar as the disagreement constraint turned out to be implausible, the current objection avoids these problems.

While an objection based on agreement avoids some of the problems with an objection based on disagreement, this is not to say that it is unproblematic. There are reasons to think that it is still not clear that the agreement must be targeting the modal proposition. For instance, consider the dialogue in (74).

- (74) a. Alex: For all I know you left them in the car.
b. Billy: You're right. Let me check.

It appears that Billy agrees with Alex, and that her response is appropriate. But if Billy's response in (74a) is appropriate, it is not surprising that Billy's response is (73a) is appropriate as well. After all, contextualists can treat (73) and (74) as having the same truth-conditions. In fact, the reasons for thinking that the agreement cannot be targeting the modal proposition expressed by (73a) are also reasons for thinking that the agreement cannot be targeting the proposition expressed by (74a). Furthermore, nothing has changed as far as the embedded proposition is concerned. We can still stipulate that Billy does not believe that the keys are in the car.

This leaves us with several options. One option is to say that the agreement is targeting the proposition expressed by (74a) after all. Another option is to say that there is a sense in which the agreement can target the embedded proposition even if Mary does not believe that proposition. A third option is to say that the agreement is targeting some other proposition. It is not clear what the right option is. But whatever the correct explanation turns out to be, there is no reason to think that it could not be extended to (73).

There are more general reasons for not wanting to rely too heavily on considerations involving agreement. Cappelen and Hawthorne (2009) have drawn attention to certain complications regarding the use of agreement reports. Suppose that Mary is a basketball coach and only applies the predicate ‘tall’ to someone who is taller than 6 feet and 8 inches. However, John is just an ordinary person and applies the predicate ‘tall’ to anyone over 6 feet. If Harry is 6 feet and 4 inches tall, it does not seem correct to report that Mary and John disagree whether John is tall. But if John is 7 feet tall, it seems appropriate to say that Mary and John agree that Harry is tall. In this case, John counts as ‘tall’ by both Mary and John’s standards. Based on these considerations, Cappelen and Hawthorne conclude that one has to be careful about relying on judgements about agreement.²⁵ Loosely speaking, one might say that agreement sometimes comes easier than disagreement. If that is true, one might be worried about relying exclusively on considerations involving agreement.

The current objection to contextualist treatments of epistemic modals like ‘might’ is based on considerations involving agreement. However, we can generate similar problems involving disagreement if we shift the focus from epistemic modals like ‘might’ to probability operators like ‘probably’ and ‘likely’.²⁶ For instance, consider the dialogue in (75) as taking place while Mary and John are looking for Harry.

- (75) a. John: Harry is probably in the office.
b. Mary: That’s not true. He might be there, but it isn’t likely.

If a contextualist treatment of ‘probably’ is correct, it is problematic to treat the disagreement as targeting the modal proposition. One would not expect Mary to be rejecting a proposition about what is probable in view of what John knows. Furthermore, it is also problematic to treat the disagreement as targeting the embedded proposition. Mary does not have to believe that Harry is not in the office in order for her response to be appropriate. In fact, it is clear that she does not take herself to be in a position to rule out that possibility that he is in the office. This suggests that this problem cannot be dealt with by interpreting the disagreement as targeting the embedded proposition.

²⁵See e.g. also Cappelen and Hawthorne (2011) and Weatherson (2011) for relevant discussion.

²⁶Thanks to Josh Dever and Brian Weatherson (p.c.) for making me appreciate this point.

I am not going to present a solution to this problem here. Doing so would most likely require a more thorough discussion of probability operators like ‘probably’ and that is beyond the scope of the present discussion. My main concern is with epistemic uses of modals like ‘might’, not with probability operators like ‘probably’. Having said that, we have seen how complicated matters involving disagreement are. We have also seen that there are cases that are hard for relativists to explain. In light of these considerations, we should not jump to the conclusion that cases like (75) call for a relativist treatment.

Conclusion

There seems to be a common assumption behind the use of both disagreement and retraction data. The underlying assumption is that the disagreement and retraction must target the modal proposition, but as we have seen, that assumption is highly suspect in both cases. I have argued that we can understand the disagreement and retraction in the relevant cases as targeting the embedded proposition rather than the modal proposition. If that is correct, it looks like we can hold on to a contextualist semantics for epistemic modals while paying sufficient respect to the disagreement and retraction data.

I have also tried to answer a potential challenge to the proposal I am defending based on the idea that disagreement and retraction can only target a proposition that is believed or asserted. When a sufficiently wide range of cases are examined, that requirement turns out to be too strong, both in the case of disagreement and in the case of retraction. Perhaps there are weaker alternatives available, but there is no obvious reason to think that they will rule out the possibility that the disagreement and retraction are targeting the embedded proposition in the cases involving epistemic modals.

As an alternative to thinking about the disagreement in terms of a difference in outright beliefs, I have suggested that we can think about the disagreement in terms of a difference in credences or degrees of belief. I have also explored issues involving the pragmatics of epistemic modals. An interesting idea is that we can think of sentences containing ‘might’ as being used to perform a speech act involving the embedded proposition in addition to an assertion of the modal proposition. It is not clear what the best way of implementing that idea is, but

using Swanson's (2006) notion of 'doxastic advice' seems promising.

In any case, this is not meant to be the final word on these matters. For instance, we have seen that there are also problems involving agreement. While there are reasons to think that those problems can be dealt with, it is not clear that we have a good understanding of exactly what is going on in the relevant cases. Moreover, I have focused almost exclusively on epistemic uses of modals like 'might'. Shifting the focus from epistemic modals like 'might' to probability operators like 'probably' introduces additional complications. Even if these problems do not call for a relativist treatment, there is certainly more work to be done.

Chapter 5

Knowledge Ascriptions and Indirect Speech Reports

Introduction

Contextualism about knowledge ascriptions remains a controversial view. Both defenders and critics of contextualism have appealed to a variety of linguistic evidence, either in an attempt to provide support for contextualism or to undermine it. The goal of the present discussion is not to settle whether contextualism about knowledge ascriptions is true, or even whether the balance of evidence is in its favour. Rather, the focus is on certain methodological issues. The overall question is what different kinds of linguistic evidence can teach us about the semantics of knowledge ascriptions. In particular, I want to discuss what we can learn by looking at the behaviour of knowledge ascriptions embedded under attitude verbs like ‘says’ or ‘believes’. A number of philosophers have argued that contextualism has problems when it comes to dealing with knowledge ascriptions embedded in indirect speech or attitude reports.¹

I will argue that there are ways for contextualists to deal with these problems without compromising the important commitments of their views. From this point of view, I will be arguing for a contextualist-friendly conclusion. But

¹See e.g. Hawthorne (2004), Cappelen and Lepore (2005), Brogaard (2008), and MacFarlane (2009).

the question is not just whether such evidence directly supports or undermines contextualism. There is a more general question of what it can teach us about the semantics of knowledge ascriptions. While considerations involving indirect speech or attitude reports do not provide any direct evidence against contextualism, it may still teach us something about the semantics of knowledge ascriptions. In order to accommodate the relevant data, defenders of contextualism may have to make further assumptions about the semantics of knowledge ascriptions. These assumptions may in turn provide the basis for further inquiries. If we look at it from this point of view, the considerations involving indirect speech or attitude reports are still relevant for the overall assessment of contextualism.

In §5.1 I say what a contextualist view about knowledge ascriptions amounts to. In §5.2 I explain why considerations involving indirect speech reports present a problem for contextualism. In §5.3 I present nonindexical contextualism as an alternative to contextualism, and explain why it does better than contextualism when it comes to dealing with indirect speech reports. In §5.4 I argue that it is not enough to say that the relevant reports are instances of mixed quotation. In §5.5 I argue that a salience-based account is also not going to explain all the data. In §5.6 I draw attention to Lewis' (1980) distinction between the semantic value of a sentence and its propositional content. In §5.7 I suggest, based on Lewis' distinction and recent work by Ninan (forthcoming), that we can deal with the problem of indirect speech reports by adding an epistemic standard parameter to the index and letting 'says' shift the value of that parameter. In §5.8 I distinguish this view from nonindexical contextualism. In §5.9 I discuss another way of dealing with indirect speech reports that makes use of the machinery of lambda abstraction. In §5.10 I point out that the accounts considered in §5.7 and §5.9 are associated with substantive assumptions regarding the semantic implementation of contextualism. §5.11 I briefly discuss the possibility of treating 'says' as a context shifting operator, a so-called 'monster', in order to avoid these commitments. In §5.12 I distinguish the problem of indirect speech reports from a problem that has to do with attributions of truth and falsity.

5.1 Contextualism about Knowledge Ascriptions

Contextualism about knowledge ascriptions is a view about the semantics of knowledge ascriptions that has been defended, in one form or another by, among others, Cohen (1988, 1999), DeRose (1992, 2009), and Lewis (1996). According to this view, knowledge ascriptions are context-dependent in the sense that they express different propositions in different contexts. Furthermore, this is not meant to include any context-dependence that is not associated with ‘knows’ or its cognates, nor does it include any context-dependence associated with features like tense.

For the purpose of the following discussion I will be talking about knowledge ascriptions being associated with different epistemic standards in different contexts.² According to contextualism, the epistemic standards that a subject must satisfy in order for a knowledge ascription to be true, varies from context to context. In one context, it might be sufficient if the subject is in a moderately strong epistemic position with respect to the relevant proposition, whereas in a different context, a much stronger epistemic position might be required.

This way of presenting contextualism about knowledge ascriptions still leaves it a fairly open question exactly how we should characterise the context-dependence of knowledge ascriptions. A natural thought is to treat ‘knows’ itself as an indexical that receives its value directly from the context. But the characterisation of contextualism is also compatible with the context-dependence being traceable to a covert element associated with ‘knows’.

I think it is safe to say that the question of how to implement contextualism semantically has received less attention than, say, questions about how contextualism can be applied to various epistemological problems.³ Cohen (1999) makes

²It is useful to talk about knowledge ascriptions as being associated with different epistemic standards in different context, but this is only in order to simplify the discussion. It is not meant to be a substantive assumption in the sense that a contextualist semantics has to make essential use of the notion of an epistemic standard. For instance, the view that Schaffer (2004) calls ‘contrastivism’ is a contextualist view in the sense I am interested in, but it is not entirely natural to characterise that view in terms of knowledge ascriptions being associated with different epistemic standards in different contexts.

³This does not amount to the claim that questions about the semantic implementation of contextualism has not received any attention at all. See e.g. Schaffer (2004) and Ludlow (2005, 2008) for relevant discussion. Stanley (2005) also pays a lot of attention to these issues in his critique of

it explicit that he is aware of the different possibilities when it comes to the semantic implementation of contextualism, but he also claims that these issues do not matter much from an epistemological point of view.

How from the view point of formal semantics should we think of this context-sensitivity of knowledge ascriptions? We could think of it as a kind of indexicality. On this way of construing the semantics, ascriptions of knowledge involve an indexical reference to standards. So the knowledge predicate will express different relations (corresponding to different standards) in different contexts.

But we could instead view the knowledge predicate as expressing the same relation in every context. On this model, we view the context as determining a standard at which the proposition involving the knowledge relation gets evaluated. So we could think of knowledge as a three-place relation between a person, a proposition, and a standard.

These semantic issues, as near as I can tell, are irrelevant to the epistemological issues. As long as we allow for contextually determined standards, it doesn't matter how formally we construe the context-sensitivity. (Cohen, 1999, p. 61)

I suspect that Cohen is right that these semantic issues are not particularly relevant to the epistemological issues. From an epistemological point of view, contextualists are warranted in setting aside these semantic issues. But insofar as we are interested in assessing the overall plausibility of the project of providing a contextualist semantics for knowledge ascriptions, we would do well in paying attention to the semantic issues. In this paper I will be arguing that these issues become relevant when we think about problems having to do with indirect speech reports. In particular, I will argue that the most promising ways of dealing with the problem of indirect speech reports require contextualists to make substantive assumptions about the semantic implementation of contextualism.

contextualism.

5.2 Indirect Speech Reports

According to contextualism, a sentence like (76) expresses different propositions in different contexts.

(76) Smith knows that he has hands.

Suppose that Mary and John are in different contexts such ‘knows’ is associated with different standards in their respective contexts. In that case, (76) might express a true proposition in Mary’s context, but a false proposition in John’s context. Suppose further that Mary sincerely utters (76), and that John, having heard what she is saying, reports what she says by uttering (77).

(77) Mary says that Smith knows that he has hands.

Given that Mary and John are in different contexts, one would think that contextualism predicts that (77) is false. By uttering (77), John is reporting that what Mary says is the proposition expressed by (76) in John’s context, not the proposition expressed by (76) in Mary’s context. But there is no reason to suppose that that is what Mary says, and we would thus expect (77) to be false. The problem is that this prediction does not seem to be correct. It appears that John can make a true report of what Mary says by uttering (77) even if they are in different contexts.

In order to appreciate the problem it might be helpful to look at an example involving the indexical ‘I’. Suppose that Mary utters (78).

(78) I am hungry.

It is clear that John cannot use (79) to report what Mary is saying.

(79) Mary says that I am hungry.

The problem is that the ‘I’ in (79) refers to John, not Mary. In order to report what Mary says, John could use the third-person pronoun ‘she’, as in (80).

(80) Mary says that she is hungry.

Alternatively, he could use direct quotation, as in (81).

(81) Mary says, ‘I am hungry’.

Both (80) and (81) are acceptable ways of reporting what Mary says, but (79) is not. If contextualism was correct, we might expect knowledge ascriptions to behave in a similar way when embedded under ‘says’, but that is not the case. John can use (77) to report what Mary says by uttering (76) even if they are in different contexts.

The challenge for contextualists is twofold. First, they need to account for the apparent truth of (77). Second, they need to explain why knowledge ascriptions do not seem to exhibit the same embedding behaviour as indexicals like ‘I’.

5.3 Nonindexical Contextualism

If contextualists were unable to deal with the problem of indirect speech reports, it is natural to think of this as providing evidence for an invariantist view of knowledge ascriptions. According to an invariantist view, knowledge ascriptions are not context-dependent in any epistemologically interesting sense. Since invariantists do not take knowledge ascriptions to express different propositions in different contexts, they do not have problems when it comes to indirect speech reports.

However, there are other alternatives to contextualism. For the purpose of the following discussion I will focus on a view that has recently been discussed by MacFarlane (2009). He has argued that the problem of indirect speech reports can be used to motivate a view he calls ‘nonindexical contextualism’.⁴ Nonindexical contextualism is interesting in this context because it promises to preserve some of the original spirit of contextualism. According to this view, knowledge ascriptions do not express different proposition in different context.⁵ Instead they express propositions that are true or false relative to some appropriate parameter. For our purposes it makes sense to follow MacFarlane in taking this to be an epistemic standard parameter.

This does indeed seem to get around the problem. On this view, (76) expresses the same proposition in Mary’s context and in John’s context.

⁴MacFarlane ultimately prefers a relativist treatment of knowledge ascriptions, but that does not matter for our current purposes. The difference between the views is that according to a relativist treatment of knowledge ascriptions, the epistemic standard parameter is determined by the context of assessment rather than the context of use. See MacFarlane (2005a, forthcoming-b).

⁵See also Brogaard (2008) who offers a similar argument for a similar view.

(76) Smith knows that he has hands.

It is not a problem for this view that John can use (77) to report what Mary says.

(77) Mary says that Smith knows that he has hands.

The original worry was that John is reporting that what Mary says is the proposition expressed by (76) in John's context, not the proposition expressed by (76) in Mary's context. But that is the same proposition on this view. In this respect, nonindexical contextualism is similar to invariantism.

This sort of view exploits the fact that the truth-value of a sentence can depend both on the context insofar as it plays a role in determining the proposition, and on the index with respect to which we evaluate the proposition for truth and falsity. In other words, it exploits the fact that a sentence is true relative to both a context and an index.⁶ Instead of treating the index as, say, just a possible world, the nonindexical contextualist treats it as pair consisting of a possible world and an epistemic standard. The standard contextualist and the nonindexical contextualist can agree on the distribution of truth-values, that is, which sentences are true at which contexts. What they disagree about is whether the contextual variation in truth-values is a result of different propositions being expressed at different contexts, or the same proposition being true or false relative to different indices.

We can give a simple semantics for 'says' and 'knows' that captures what the nonindexical contextualists are after. This will be useful later on when we compare nonindexical contextualism to other views. In order to do that, we need to relativise the denotation $\llbracket \phi \rrbracket$ of an expression ϕ to a context c , an epistemic standard n , and a possible world w . For the purpose of the following discussion it will be convenient to assume that the context determines an assignment of values to variables and therefore I will not add a separate parameter for the variable assignment. Moreover, in order to label the types of denotation, we can take e to be the type of individual, t to be the type of truth-values, s to be the type of possible worlds, and j to be the type of epistemic standards. This allows us to give a semantics for 'knows' and 'says' along the following lines:⁷

⁶See e.g. Lewis (1980) and Kaplan (1989) for a more thorough discussion of the motivations behind this kind of semantic framework.

⁷The semantics is based on the semantics for attitude verbs given by Heim and Kratzer (1998,

(82) $\llbracket \text{says} \rrbracket^{c,n,w} = \lambda p_{\langle j, \langle s, t \rangle \rangle} \cdot \lambda x_e$. For every world w' and epistemic standard n' compatible with what x says in $w : p(n')(w') = 1$

(83) $\llbracket \text{knows} \rrbracket^{c,n,w} = \lambda p_{\langle j, \langle s, t \rangle \rangle} \cdot \lambda x_e$. For every world w' and epistemic standard n' compatible with what x knows relative to n in $w : p(n')(w') = 1$ ⁸

This semantics makes ‘knows’, but not ‘says’, sensitive to the epistemic standard parameter of the index. A sentence like (76) has different truth-values relative to different epistemic standards. However, it does not follow that a sentence like (77) has different truth-values relative to different epistemic standards. Instead of involving quantifying over just worlds, attitude verbs like ‘says’ and ‘knows’ quantify over worlds and epistemic standards. Simply put, the thought is that (77) is true just in case the proposition that Smith knows that he has hands is true relative to every world-standard pair compatible with what Mary says. Since we are quantifying over epistemic standards, the truth-value of (77) does not depend on the epistemic standards that are relevant in John’s context. This reflects the idea that a proposition, what a speaker says, is true or false relative not only to a world, but a world and an epistemic standard, and that (77) is true just in case what Mary says is the proposition that Smith knows that he has hands.

If contextualists do not have a satisfactory answer to the problem of indirect speech reports, this could be seen as providing motivation for nonindexical contextualism. However, while I agree that the problem of indirect speech reports requires contextualists to make substantive assumptions about the semantics of knowledge ascriptions, I do not think we need to go as far as to adopt a version of nonindexical contextualism. In the next two sections I will look at two attempts at dealing with the problem of indirect speech reports that do not require any substantive assumptions about the semantics of knowledge ascriptions. I will argue that these proposals fail to deliver a solution to the problem. Having done that, I will look at other proposals that are more promising, but which require more substantive semantic assumptions.

ch. 12) and von Fintel and Heim (ms.).

⁸For the purpose of the following discussion, I am ignoring some difficult questions about to handle the factivity of ‘knows’. See e.g. Stanley (2005) and Brogaard (2008) for relevant discussion.

5.4 Mixed Quotation

A proposal that has been put forward by Hawthorne (2006) and Cappelen and Hawthorne (2009) is that the relevant examples could involve what Cappelen and Lepore (1997) have called ‘mixed quotation’. That means that part of the sentence is quoted as in (84).

(84) He now plans to make a new, more powerful absinthe that he says will have ‘a more elegant, refined taste than the one I’m making now’.⁹

What is important for our purposes, is that the part of the sentence that is quoted can contain context-dependent expressions. If there were no quotation marks in (84), the indexical ‘I’ would refer to the person making the report, but that is not the case when the relevant part of the sentence is quoted. The question is whether we can treat (77) as involving mixed quotation in the same way as (84).

(77) Mary says that Smith knows that he has hands.

If the relevant parts of the sentence are quoted in this way, it is not surprising that John can use (77) to report what Mary says. After all, there is nothing problematic about John using (85) to report what Mary says.

(85) Mary says, ‘Smith knows that he has hands’.

An immediate and obvious worry with this proposal is that there are no visible quotation marks in (77). However, it is not clear that this is a decisive objection to this line of response. Cappelen and Hawthorne argue that we should not expect people to be careful when it comes to using quotation marks. As Hawthorne (2006, p. 446) puts it, ‘there are no use-mention police in the ordinary world’. Cappelen and Hawthorne (2009, p. 59) observe that it is natural to understand (86) as involving a mixed quotational use of ‘gay’ if A meant ‘happy’ and B meant ‘homosexual’.

(86) A and B both said that they are gay, but they mean very different things by that.

⁹This example is from Cumming (2005).

Having granted this point, it still needs to be said that these matters are somewhat delicate. In the cases involving indexicals like 'I' it is much harder to hear the relevant reports as involving mixed quotation unless quotation marks, or some other way of indicating that a part of the sentence is quoted, are used. For instance, consider the difference between (84) with and without the quotation marks. In the latter case, it is natural to hear the 'I' as referring to the person making the report.

In any case, it is not clear that appealing mixed quotation is enough to get contextualism completely off the hook. There is a further worry which is that the mixed quotation strategy does not explain why 'knows' does not behave in the same way as indexicals like 'I' when embedded in indirect speech reports. If it is relatively easy to hear (77) as involving mixed quotation, why is it much harder to hear (79) as involving mixed quotation?

(79) Mary says that I am hungry.

It would be surprising if it was easy to have mixed quotation when the relevant part of the sentence contains 'knows', but not when 'I' is involved. After all, there are cases of mixed quotation involving indexicals, as illustrated by examples like (84). In order for there to be such a difference, the possibility of a report involving mixed quotation would have to be dependent on whether the relevant part of sentence contains 'knows' or 'I'. However, it is far from obvious that this should make a difference as far as mixed quotation is concerned. Why should there be such restrictions on mixed quotation? This makes it hard to see how mixed quotation could provide a complete solution to the problem of indirect speech reports. Anyone who tries to appeal to mixed quotation in order to explain why it is easy to use to (77) to report what Mary is saying, faces the objection that this account also predicts that it should be just as easy in the case of (79).

There is certainly a case to be made that we need to be aware of the possibility that indirect speech reports may involve mixed quotation. In some cases, there might even be strong pressure to think that mixed quotation is involved even if the report does not contain any visible quotation marks. However, that does not mean that mixed quotation can fully explain the fairly robust patterns of data involving indirect speech reports and knowledge ascriptions.

5.5 Salience and Parasites

A natural response on behalf of contextualism is to say that when John is reporting what Mary says, what is salient is what is going on in Mary's context. The standards associated with 'knows' in (77) are the standards in Mary's context.

(77) Mary says that Smith knows that he has hands.

According to the terminology of Cappelen and Hawthorne (2009), John's use of 'knows' is 'parasitic' on the relevant features of Mary's context. According to this proposal, the complement clause in (77) expresses the same proposition as in Mary's context because it is the relevant features of her context that are relevant. This explains why we can hear (77) as true even though Mary and John are in different contexts.

A response along these lines has been suggested by Hawthorne (2006) and Humberstone (2006), with Humberstone basing his proposal on certain observations made by Nunberg (1993). However, while they discuss problems involving indirect speech reports, they have not primarily been concerned with knowledge ascription. Instead, a lot of the discussion has focused on expressions like 'local' and 'nearby'. These expressions seem to require a contextually determined location, but they appear to be flexible in terms of what the contextually relevant location is. More specifically, the relevant location does not have to be the location of the speaker. As Cappelen and Hawthorne (2009) have pointed out, someone uttering (87) could be talking about Naomi going to a beach nearby to their own location, but they could also be talking about Naomi going to a beach nearby to some other salient location.

(87) Naomi went to a nearby beach.

To see why this is relevant, suppose that someone overhears an utterance of (87) by Nicole. It appears that they can use (88) to report what Nicole says even if they are not in the same location as Nicole.

(88) Nicole says that Naomi went to a nearby beach.

However, this is perhaps not very surprising given the flexibility of 'nearby'. It could be that the use of 'nearby' in (88) is parasitic on the location that is relevant

in Nicole's context. If this is right, (88) will come out as true even if the speaker is not in the same location as Nicole.

There is also a story to be told about why 'nearby' and 'I' seem to behave differently when embedded in indirect speech reports. The difference between 'nearby' and 'I' is that in the case of 'nearby' there is a certain amount of flexibility when it comes to determining the relevant location. It does not have to be the location of the speaker. On other hand, it is much harder to get 'I' to refer to someone other than the speaker. That is why we can typically only interpret 'I' as referring to John and not Mary in (79).

(79) Mary says that I am hungry.

However, there are still several worries about this line of response. Even if this is a good explanation of what is going on in the case of expressions like 'local' and 'nearby' it is not clear that this line of response can be extended to knowledge ascriptions. While it is harder to recognise different readings in the case of knowledge ascriptions, MacFarlane (2009) has argued that there are reasons for thinking that knowledge ascriptions are less flexible than expressions like 'local' and 'nearby' in the relevant sense. As Stanley (2005, p. 67) has observed, utterances of sentences like (89a) and (89b) are generally held to be infelicitous.

- (89) a. Bill knows that he has hands, but Bill does not know that he is not a bodiless brain in a vat.
b. Bill does not know that he is not a bodiless brain in a vat, but Bill knows he has hands.

This would be surprising if the different occurrences of 'knows' could be associated with different standards.¹⁰ In that case, it should be possible to get a true reading of sentences likes this, but such a reading does not seem to be available. According to MacFarlane, this makes it more difficult for contextualists to respond to the problem of indirect speech reports. The explanation of why (88) can be true as uttered by someone who is not in the same location as Nicole exploits the flexibility of 'nearby' when it comes to determining the relevant location. The worry is that knowledge ascriptions are not flexible to the same extent. MacFarlane argues that contextualists face a dilemma. If they want to explain

¹⁰For further relevant discussion see e.g. Ichikawa (forthcoming).

the apparent truth of (77) in this way, they must grant that sentences like (8ga) and (8gb) can be true.

A worry with this argument is that there are several ways in which an expression can be flexible. What examples like (8g) show is that the epistemic standards associated with ‘knows’ do not easily shift in such a way as to allow the different occurrences of ‘knows’ to be associated with different standards. However, what is at issue here is whether knowledge ascriptions are flexible with respect to what the relevant standards are and how they are determined. If these things can come apart, it is not clear how effective MacFarlane’s argument is.

In any case, it is not clear how different ‘knows’ is from ‘nearby’ with regards to acceptability of sentences like (8ga) and (8gb). Insofar as ‘nearby’ is fairly flexible with respect to what the relevant location can be, we might expect (90) to have a true reading on which the different occurrences of ‘nearby’ are associated with different locations.

(90) Naomi took the bus to a nearby beach, but she did not go to a nearby beach.

However, an utterance of (90) would typically be just as infelicitous as an utterance of (8ga) or (8gb).¹¹ In other words, it appears that MacFarlane’s reasons for thinking that ‘knows’ is different from ‘nearby’ are not all that compelling.

In any case, there is a more straightforward reason to doubt that a salience-based account of speech reports is adequate. Even if it captures what is going on in examples like (77), it fails to provide an account that can be extended to all indirect speech reports. Cappelen and Hawthorne (2009) argue that it does not explain the apparent correctness of collective speech reports like (91).

(91) Mary and John say that Smith knows that he has hands.

Even if Mary and John were in different contexts when they uttered (76), it appears that Harry can use (91) to report what they are saying. But in this case we cannot say that it is what is going on in the subject’s context that is relevant, since Mary and John are in different contexts. The relevant standards cannot be either the standards associated with ‘knows’ in Mary’s context or the standards associated with ‘knows’ in John’s context. In either case, we would get the wrong

¹¹Thanks to Andreas Stokke for pointing this out to me.

result as the standards would either fail to match the standards in Mary's context or the standards in John's context.

It therefore appears that the line of response under discussion does not do very well when it comes to dealing with collective speech reports. Insofar as the data seems to be just as problematic in the case of collective speech reports, contextualists need to find a way of dealing with such reports. Having said that, there still seems to be something right about the idea that we want the relevant standards to be the standards of the subject of the report. What examples like (91) show, is that such a story needs to be implemented differently or supplemented with a different story that can handle collective reports.

5.6 Semantics and Propositional Content

The proposals we have considered so far have not forced us to make any major assumptions about the semantics of knowledge ascriptions or indirect speech reports. While I do not claim that I have surveyed every possible way of explaining the data without making significant assumptions about the semantics of knowledge ascriptions or indirect speech reports, it makes sense to see if we can do better if we are willing to make such assumptions.

In order to see what kind of assumptions we have to make, we need to pay attention to a distinction drawn by Lewis (1980). He pointed out that we can distinguish between the semantic value of a sentence and its propositional content. The former is assigned as a part of providing a compositional semantics for the language, whereas the latter is meant to play the role as the objects of attitudes and speech acts. Lewis argued that the semantic value and the propositional content of a sentence may come apart. In order to provide a compositional semantics for a language that contains modal and temporal operators, we can take indices to have a world and a time parameter that can be shifted by the relevant operators. In that case we may think of the semantic value of a sentence as being represented by a function from indices, world-time pairs in this case, to truth-values.¹²

¹²As Lewis (1980) pointed out it does not matter whether we think of sentences as having variable, but simple semantic values or constant, but complicated semantic values. If it is the former, semantic values can be represented simply as functions from indices to truth-values, but sentences

But this does not require us to treat the propositional content, the thing that we assert, believe, doubt, and so forth, in the same way. It is compatible with representing the propositional content of a sentence as a function from worlds to truth-values. Lewis observed that we can define the propositional content of a sentence ϕ in a context c as the proposition that is true at a world w if and only if ϕ is true at context c and the index that results if we take the index i_c of context c and shift its world parameter to w .¹³

In light of these considerations, there is a question about how to understand the difference between standard contextualism and nonindexical contextualism. Should nonindexical contextualism be understood as a view about the semantic value or the propositional content of knowledge ascriptions? Insofar as MacFarlane does not distinguish between the semantic value and the propositional content of a sentence it is natural to think that he does not take them to come apart in this case. In any case, this raises the possibility that one can take indices to contain an epistemic standard parameter without being committed to the view that the propositional content of a sentence is true or false only relative to an epistemic standard. This raises the further question of whether we need to go as far as to adopt a nonindexical contextualist semantics for knowledge ascriptions in order to deal with indirect speech reports, or whether we can exploit Lewis' distinction while preserving more of the spirit of the original contextualist proposal. In the next section I will look at one proposal which promises to do just that.

5.7 The Operator Account

One way of trying to solve the problem of indirect speech reports is to argue that indices contain an epistemic standard parameter and that 'says' is an operator that shifts this parameter. This sort of response has recently been explored by Ninan (forthcoming) in the case of epistemic modals. He is responding to a challenge much the like the one facing contextualist treatments of knowledge

may have different semantic values in different contexts. If it is the latter, semantic values are represented as functions from contexts and indices to truth-values.

¹³Lewis (1980) made it clear that this is not the only way in which one can define the propositional content of a sentence in a context. We can for instance also define something like the diagonal propositional in Stalnaker's (1978) sense.

ascriptions. The worry is that contextualist treatments of epistemic modals cannot explain the behaviour of epistemic modals that are embedded in attitude reports.¹⁴ Ninan’s solution is to treat attitude verbs like ‘thinks’ as operators that shift the relevant parameter of the index. In the case of epistemic modals he simply takes this to be an individual parameter, and the attitude verb shifts the value of this parameter to the subject of the report. In other words, if the subject of the report is Sam, the value of the individual parameter is shifted to Sam.

However, Ninan points out that this does not mean that we have to treat the propositional content as true or false relative to different individuals. This is just an application of the point made by Lewis (1980) regarding the relationship between the semantic value of a sentence and its propositional content. We can still represent the propositional content as a function from worlds to truth-value, even if we represent the semantic value of a sentence as, say, a function from world-time-individual triples to truth-values. Just because the index contains an individual parameter it does not have to be the case that the propositional content varies in truth-value across individuals.

A similar strategy is possible in the case of knowledge ascriptions. In this case, we have an epistemic standard parameter instead of an individual parameter, but it is still feasible to treat attitude verbs like ‘thinks’ or ‘says’ as shifting the value of this parameter to the standards of the subject of the report. If the subject of the report is Mary, the value of the epistemic standard parameter is shifted to Mary’s standards. Let us call this the ‘operator account’. This is a view about the semantics of both knowledge ascriptions and indirect speech reports.

In order to state the view more formally, we need to relativise the denotation $\llbracket \phi \rrbracket$ of an expression ϕ to a context c , an epistemic standard n , a time t , and a possible world w . The only reason for including a time parameter is that it allows us to talk about the epistemic standards of an individual at a time and in a world. Let $n_{x,t,w}$ be the epistemic standards of an individual x at a time t in a world w . The value of the time parameter will not be shifted by ‘says’ or ‘knows’, and I am not going to discuss the interaction between attitude verbs and tense. When it comes to the semantic types, we can take e to be the type of individual, t to be the type of truth-values, s to be the type of possible worlds, i to be the

¹⁴See Egan et al. (2005) and Weatherson (2008) for a more detailed discussion of this sort of challenge for contextualist treatments of epistemic modals.

type of time, and j to be the type of epistemic standards. This allows us to give a semantics for ‘says’ along the following lines:

$$(92) \quad \llbracket \text{says} \rrbracket^{c,n,t,w} = \lambda p_{\langle j, \langle i, \langle s, t \rangle \rangle \rangle} \cdot \lambda x_e. \text{ For every world } w' \text{ compatible with what } x \text{ says at } t \text{ in } w : p(n_{x,t,w})(t)(w') = 1$$

This is meant to capture the idea that ‘says’ shifts the value of the epistemic standard parameter to the epistemic standards $n_{x,t,w}$ of individual x at time t in world w . Since ‘knows’ is an attitude verb, we presumably want it to be able to shift the epistemic standard parameter in the same way. We also want it to be the case that the truth-value of a knowledge ascription is dependent on the value of the epistemic standard parameter. The following semantics for ‘knows’ delivers both these results:

$$(93) \quad \llbracket \text{knows} \rrbracket^{c,n,t,w} = \lambda p_{\langle j, \langle i, \langle s, t \rangle \rangle \rangle} \cdot \lambda x_e. \text{ For every world } w' \text{ compatible with what } x \text{ knows at } t \text{ in } w \text{ relative to } n : p(n_{x,t,w})(t)(w') = 1^{15}$$

This view seems to deliver the right result in the case in which John is using (77) to report what Mary is saying when she utters (76).

(76) Smith knows that he has hands.

(77) Mary says that Smith knows that he has hands.

The original worry was that by uttering (77), John is reporting that what Mary says is the proposition expressed by (76) in John’s context, not the proposition expressed by (76) in Mary’s context. However, according to the view under discussion, ‘says’ shifts the value of the epistemic standard parameter to Mary’s standards. That means that (77) is true just in case all the worlds compatible with what Mary says are worlds in which Smith knows that he has hands relative to Mary’s epistemic standards. In other words, it is true just in case Mary stands in the relation denoted by ‘says’ to the proposition that Smith knows that he has hands relative to Mary’s epistemic standards. This seems to be what we want to say in this case.

The view also seems to have the resources to deal with collective reports like (91).

¹⁵I am still ignoring issues concerning the factivity of ‘knows’.

(91) Mary and John say that Smith knows that he has hands.

The idea is that (91) is true at a context c and an index $\langle n, t, w \rangle$ just in case both Mary and John have the property picked out by the verb phrase ‘says that Smith knows that he has hands’.

(94) $\llbracket \text{says that Smith knows that he has hands} \rrbracket^{c,n,t,w} = \lambda x_e$. For every world w' compatible with what x says at t in w : $\llbracket \text{Smith knows that he has hands} \rrbracket^{c,n_x,t,w,t,w'} = 1$

Roughly speaking, for both Mary and John to have this property is for each of them to have the property of being an x such that x says that Smith knows that he has hands relative to x 's standards. This seems to give us the kind of truth-conditions that we want.

If something like this account is on the right track, we have in effect improved on the idea that the relevant epistemic standards are the standards of the subject. Instead of understanding this in terms of the standards of the subject being salient, we are treating ‘says’ as an operator that shifts the value of the epistemic standard parameter to the standards of the subject. Still, it is worth noting that in order to do this we have had to make certain assumptions about the semantics of knowledge ascriptions as well as the semantics of indirect speech reports. On the proposed semantics for ‘knows’, it is sensitive to the epistemic standard parameter of the index. In this sense, it does not behave like an indexical that denotes different relations between an individual and a proposition in different contexts. At this point I am not making any claim as to whether that is a welcome or an unwelcome consequence. The point is that it matters how we think about the context-dependence of knowledge ascriptions.

5.8 Nonindexical Contextualism After All?

At this point one might wonder whether the operator account just amounts to a version of nonindexical contextualism. This looks suspiciously like a terminological question that is not particularly philosophically interesting in its own right, but I still think there is something to be said for thinking that the operator account differs from the kind of nonindexical contextualist view that MacFarlane

(2009) is talking about, in certain important respects. According to the operator account, there is a good sense in which we can still think of knowledge ascriptions as expressing different propositions in different contexts, insofar as the propositional content varies from context to context. This is because, as Lewis (1980) pointed out and as Ninan (forthcoming) has recently emphasised, we are not compelled to identify the propositional content of a sentence with its semantic value. If we adopt (95) as a definition of propositional content, knowledge ascriptions will still have different propositional content in different contexts.

- (95) The propositional content of a sentence ϕ in a context c is the proposition that is true at a world w if and only if ϕ is true at context c and the index that results if we take the index i_c of context c and shift its world parameter to w .

A reason for thinking that the operator account looks similar to nonindexical contextualism is that it does not treat knowledge ascriptions as having different semantic values in different contexts. However, it is not clear that this is something we should attach a lot of importance to. Lewis (1980) observed that it does not really make a difference if we think of the semantic value of a sentence as a function from indices to truth-values, with the sentence having different semantic values in different context, or simply as a function from context-index pairs to truth-values. In the latter case, even sentences containing indexicals like 'I' would not have different semantic values in different contexts.

Another sense in which the view is similar to nonindexical contextualism is that the truth-value of a knowledge ascription does not depend directly on the context, but on the epistemic standard parameter of the index. However, the reason for having an epistemic standard parameter in the index is to make it possible for it to be shifted by operators like 'says'. Again, it is not clear why a contextualist would attach a lot of significance to this.

Furthermore, I want to point out that the operator account does not face the same problems as nonindexical contextualism. Insofar as two views face different problems, that is at least some reason to distinguish between them. For instance, a nonindexical contextualist who takes the propositional content of a sentence to vary in truth-value across epistemic standards has to defend the view that the relevant propositional contents are suitable as the objects of attitudes and speech

acts. One might worry that it does not really make sense for the things that we believe to vary in truth-value across epistemic standards, or one might worry that it requires us to revise our picture of how assertion works. MacFarlane (2009, p. 243) discusses and attempts to address some concerns along these lines, but the point is not so much whether these worries can be addressed or not. The point is rather that a proponent of the operator account does not have to worry about these problems in the first place.¹⁶ The reason is simply that on this view the propositional content does not vary in truth-value across epistemic standards. This does not mean that the operator account does not face any problems, it only means that it does not face exactly the same problems as nonindexical contextualism.

While I want to avoid a potentially fruitless terminological discussion, these considerations seem to support the idea that there is a sense in which the view under discussion is interestingly different from the kind of nonindexical contextualist view that MacFarlane is talking about. If that is right, there appears to be sense in which considerations involving indirect speech reports do not give us a reason for giving up contextualism in favour of nonindexical contextualism. Given certain assumptions about the semantics of knowledge ascriptions and indirect speech reports, contextualists seem to have to resources to deal with the problem of indirect speech reports.

5.9 Lambda Abstraction

Once we have a proposal for dealing with indirect speech reports on the table, it becomes natural to see whether there are similar proposals that can also get the job done. In particular, one might be interested in whether there ways of doing this without treating ‘knows’ as sensitive to the index rather than the context. In this section, I will look at one such proposal. Cappelen and Hawthorne (2009), following a suggestion by Stanley (2005), have advocated using the machinery of lambda abstraction to make sense of otherwise problematic indirect speech reports. I will call this the ‘lambda abstraction account’. While this strategy avoids some of the commitments of the operator account, it will become clear that it is

¹⁶I am ignoring the possibility of a more general challenge directed against the view that we can represent the objects of attitudes and speech acts as functions from worlds to truth-values.

associated with certain other commitments.

A simple way to think of lambda abstraction is that it allows us to form an expression that denotes a property. For instance, the expression ‘ $\lambda x (x \text{ is red})$ ’ denotes the property of being red.¹⁷ This machinery already has other applications in semantics. Cappelen and Hawthorne point out that lambda abstraction is used to account for the different readings of sentences like (96).

(96) John loves his mother and Bill does too.

A sentence like (96) has two readings. On the so-called ‘strict’ reading, it is true just in case John loves John’s mother and Bill loves John’s mother too. On the so-called ‘sloppy’ reading, it is true just in case John loves John’s mother and Bill loves Bill’s mother. It is the second reading that is interesting for our purposes. In order to capture this reading, we can treat the pronoun ‘his’ as a variable that is bound by a lambda operator. The relevant reading of the verb phrase ‘loves his mother’ is then given by (97).

(97) $\lambda x. x \text{ loves } x\text{'s mother}$

We can think of (97) as denoting the property of being an x such that x loves x ’s mother. Simply put, (96) is true on a so-called ‘sloppy reading’ just in case both John and Bill have this property.

Cappelen and Hawthorne argue that the same machinery can be used to deal with the problems involving indirect speech reports. While their discussion focuses on expressions like ‘local’ and ‘nearby’, the basic strategy can be extended to knowledge ascriptions in a relatively straightforward manner. The idea is to let ‘knows’ be associated with a variable that can be bound by a lambda operator. For instance, we can think of ‘knows’ as being associated with a variable of the form $f(x)$, such that x is a variable over individuals and f is a variable over functions from individuals to epistemic standards.¹⁸ When these variables are not bound, they are assigned a value by the contextually determined assignment function. When x is not bound we may assume that the value of x will typically, but perhaps not invariably, be the speaker. There is probably a lot to be said about the

¹⁷I am using the same notation as Cappelen and Hawthorne (2009).

¹⁸The idea of treating knows as associated with a variable of the form $f(x)$ is taken from Stanley and Szabó (2000).

function that is assigned to f , but for our purposes it is convenient to think of it as a function that takes us from an individual to the epistemic standards that are relevant or appropriate to the individual in question.

What matters for our purposes is that by positing a variable that can be bound by a lambda operator, we gain additional resources when it comes to dealing with indirect speech reports. Since Cappelen and Hawthorne focus on lambda abstraction as a means of dealing with collective reports, it makes sense to start by looking at examples like (91).

(91) Mary and John say that Smith knows that he has hands.

Instead of having the context assign a value to the variable x , we posit a lambda operator that binds it. Roughly speaking, the relevant reading is given by (98).

(98) Mary and John λx (x says that Smith knows by $f(x)$ that he has hands).

On this reading, (91) is true just in case both Mary and John have the property of being an x such that x says that Smith knows by $f(x)$ that he has hands. If we assume that f is assigned a function that takes us from an individual to the standards of that individual, this seems to get us the right predictions. It is enough for (91) to be true that Mary says that Smith knows that he has hands relative to Mary's standards, and that John says that Smith knows that he has hands relative to John's standards.

It is worth noting that Cappelen and Hawthorne talk about lambda abstraction as a means to deal with the problems involving collective reports. When it comes to non-collective reports, reports like (77), they seem to favour the kind of salience-based account that was considered and rejected in §5.5.

(77) Mary says that Smith knows that he has hands.

Insofar as the main reason for rejecting salience-based accounts was that they could not deal with collective reports, a salience-based account of non-collective reports might still be viable. Still, there is also the possibility of extending the lambda abstraction strategy to non-collective reports. In that case, the relevant reading of (77) would be (99).

(99) Mary λx (x says that Smith knows by $f(x)$ that he has hands).

On this reading, (77) is true just in case Mary has the property of being an x such that x says that Smith knows by $f(x)$ that he has hands. These predictions are similar to the predictions made by the operator account. The relevant standards are the standards of the subject of the report.

The lambda abstraction account and the operator account are similar in many respects. They basically account for collective reports in the same way. What the operator account does by having ‘says’ shift the epistemic standard parameter of the index, the lambda abstraction account does by using the machinery of lambda abstraction. Simplifying somewhat, we can think of both views as taking (91) to be true just in case both Mary and John have the property of being an x such that x says that Smith knows that he has hands relative to x ’s standards.

However, the views differ in terms of the assumptions they make about knowledge ascriptions. If we compare the lambda abstraction account to the operator account, the former gives us a more traditional contextualist approach to knowledge ascriptions. Instead of treating ‘knows’ as sensitive to the epistemic standard parameter of the index, the epistemic standards are supplied by the context. The epistemic standards are determined by the values assigned to f and x by the context, and the standards will typically be the standards of the speaker. The semantics for ‘knows’ could be given by something along the lines of (100).

$$(100) \quad \llbracket \text{knows} \rrbracket^{c,w} = \lambda p_{\langle s,t \rangle}. \lambda x_e. \lambda y_e. \text{For every world } w' \text{ compatible with what } x \text{ knows in } w \text{ relative to } f(y) : p(w') = 1$$

If we adopt this semantics for ‘knows’, there is no need to relativise the denotation to an epistemic standard. The context-dependence of ‘knows’ is handled by the assignment of values to variables that is determined by the context.

A potentially relevant difference between the views is that the lambda abstraction account is more flexible than the operator account. This is a good thing if the additional flexibility is exploited in some way. However, it might be a less attractive feature of the proposal if it turns out that the additional flexibility is not doing any work. If the operator account is correct, the value of the epistemic standard parameter is always shifted to the standards of the subject. The shift is obligatory. However, the lambda abstraction account allows for the possibility that there is a reading of the relevant reports on which the variable associated with ‘knows’ is not bound.

This means that if we account for the truth of (77) or (91) in terms of lambda abstraction, we also predict that these sentences have a false reading. This is a reading on which the variable x is not bound by the lambda operator, but instead receives its value in context. After all, if a sentence like (96) also has a so-called ‘strict’ reading, and the lambda abstraction account is correct, we should expect the same to be true of (77) and (91). Insofar as Mary and John are in contexts in which different standards are in play, we should expect (77) and (91) to be false on this reading.

The question is whether a false reading of (77) and (91) is available. While it is often easy to report what speakers say in this way, a case can be made that there are cases in which there is something wrong with a report like (77). DeRose (2006, 2009) discusses an example like this. In his example, Thelma and Louise are at a tavern where they are discussing whether Jim was at the office yesterday. The goal of the discussion is merely to decide who has to pay up a two dollar bet on whether Jim was at the office the day before. On the basis of testimony from a reliable source, as well as having seen Jim’s hat in the hall, Louise sincerely utters (101).

(101) I know that Jim was in the office.

However, after leaving the tavern, Thelma is contacted by the police. They are investigating a serious crime, and they are interested in Jim’s whereabouts the day before. Louise has the same evidence as Thelma, she saw his hat in the hall and she has testimony from a reliable source. But while she is being interviewed by the police, this does not appear sufficient for her to ascribe knowledge to herself. In this context, an utterance of (102) by Thelma would seem appropriate.

(102) I don’t know whether Jim was in the office.

What matters for our purposes, is that there might also appear to be something wrong about her to using (103) to report what Louise said.

(103) Louise said that she knew that Jim was in the office.

If this is evidence that the relevant speech report has a genuinely false reading, this would be good news for the lambda abstraction account and bad news for the operator account. The lambda abstraction account predicts that such a reading

should be available, whereas the operator account predicts that (103) should be true as uttered by Louise.¹⁹ However, we should not be too quick to conclude that (103) has a false reading. While there might be something awkward about (103) in this context, (104) appears to be straightforwardly false.

(104) Louise didn't say that she knew that Jim was in the office.

If (103) was false, we should expect (104) to be true. This is some evidence that the inappropriateness of (103) is not due to its falsity. Instead, it could be that an utterance of (103) by Thelma is merely misleading. For instance, it might convey that Louise has more evidence, or at least took herself to have more evidence, than Thelma. Such an account would have to be backed up by a more detailed pragmatic story, and that is not a trivial task. But the point still stands that the evidence that (103) has a false reading, is pretty shaky. Insofar as we cannot detect the relevant reading, it is problematic for the lambda abstraction account to predict this kind of ambiguity. This is not in any way a decisive objection, but it would be nice to have an explanation of why the 'strict' reading is not available. It also means that DeRose's example does not necessarily provide any evidence against the operator account.

5.10 Taking Stock and Looking Ahead

Both the operator account and the Lambda lambda abstraction account deliver the right predictions when it comes to knowledge ascriptions that are embedded in indirect speech reports. However, both accounts require contextualists to make substantive assumptions about the semantics of knowledge ascriptions. From the point of view of dealing with indirect speech reports, not just any semantic implementation of contextualism will do. More specifically, neither account treats 'knows' itself as an indexical in the sense that it receives its value directly from the context. According to the operator account, 'knows' is sensitive to the index, and the context only plays a role insofar as it determines the relevant index. According to the lambda abstraction account, 'knows' is not itself an indexical. Instead, it is associated with a covert variable that is assigned a value by the context.

¹⁹I am ignoring any complications having to do with the use of past tense.

This does not have to be construed as a problem for contextualism. It was never clear that the best way of implementing contextualism was to treat ‘knows’ as an indexical.²⁰ Schaffer (2004) has offered independent reasons for thinking that ‘knows’ is not an indexical that are not dependent on considerations involving indirect speech reports. What is interesting, is whether the assumptions made by the operator account and the lambda abstraction account can be subject to further empirical testing. A detailed examination of these issues are beyond the scope of the present discussion, but I want to highlight certain avenues for future research.

In the case of the operator account, there are at least two sources of independent evidence one might look for. The first source of independent evidence would be evidence that ‘says’ can shift other parameters in the same way as the epistemic standard parameter. This would require looking at other expressions than ‘knows’ and seeing how they behave when they are embedded under ‘says’.²¹

The second source of independent evidence would be evidence that the epistemic standard parameter can be shifted by other operators. A natural candidate would be expressions like ‘by the standards of x ’.²² However, it is not clear that we want to treat these expressions in this way. As Stanley (2005, p. 70) has pointed out, expressions like this can also occur in sentences like (105) that do not contain ‘knows’ or any other epistemic terms.²³

(105) By the standards of chemistry, what is in the Hudson River isn’t water.

In this case it would be strange to think that what the expression ‘by the standards of chemistry’ is doing, is to shift the value of the epistemic standard parameter. Insofar as these expressions have a uniform semantics, this is evidence that they

²⁰See e.g. Schaffer (2004) and Ludlow (2005, 2008).

²¹It might be particularly interesting to compare knowledge ascriptions to epistemic modals. See e.g. Ninan (forthcoming) for relevant discussion.

²²Ludlow (2005) talks about expressions like these, but he does not treat them as operators that shift a parameter of the index. They simply make the relevant epistemic standards explicit.

²³It is worth mentioning that Stanley’s (2005) arguments are directed against the view defended by Ludlow (2005) to the effect that the context-dependence of ‘knows’ is traceable to a covert element that is associated with ‘knows’. However, similar considerations can also be used to cast doubt on the claim that expressions like ‘by the standards of x ’ are operators that shift the value of the epistemic standard parameter of the index.

are not operators that shift the epistemic standard parameter of the index. If that is correct, we would have to look elsewhere for independent evidence of an epistemic standard parameter that can be shifted by operators.

In the case of the lambda abstraction account, the key assumption is that ‘knows’ is associated with a variable that can be bound by a lambda operator. This is most naturally construed as a syntactic claim about sentences containing ‘knows’. If this claim is correct, it would be surprising if this variable could not be bound by an overt quantifier like ‘everyone’. In other words, we should expect a sentence like (106) to have a so-called ‘bound’ reading.²⁴

(106) Everyone knows that Smith has hands.

On a bound reading, (106) is true just in case every x knows by x ’s standards that Smith has hands. The existence of such a reading is controversial.²⁵ It is difficult to simply hear a bound reading of (106). Indeed, it might be difficult to detect such a reading directly. Instead one might want to look for a situation in which (106) is true on a bound reading, but false on a normal reading. In any case, this is not the time and the place to settle this question. For the purpose of the present discussion, it is sufficient to note that this is the kind of evidence one might want to look for.²⁶

The point of this is to show that the kinds of assumptions that are made by the operator account and the lambda abstraction account are themselves subject to further empirical testing. It might turn out that these assumptions are more or less plausible. That is something that will have to be addressed by further research. In any case, this suggests that a proper evaluation of contextualism will have to take into account a wide range of data. We started out by considering how knowledge ascriptions behave when embedded in indirect speech reports,

²⁴There has been a lot of discussion about whether a bound reading should be taken as evidence for the existence of a covert variable. See e.g. Stanley (2000). I am more interested in the claim that if we posit a covert variable, we should expect to find a reading on which it is bound. This may still not be entirely uncontroversial. See e.g. Ludlow (2005, 2008).

²⁵Schaffer (2004) argues that there is evidence of bound readings involving knowledge ascriptions. His main example involves quantification over tests or questions rather than individuals, but it might be possible to produce a variant of his example that involves quantification over individuals.

²⁶There are other considerations than binding that might be relevant. See e.g. Schaffer (2004, forthcoming-a) for other kinds of syntactic evidence that might be relevant.

and we ended up looking at how knowledge ascriptions interact with quantifiers. From the point of view of providing an overall assessment of contextualism, it matters how the view is implemented linguistically.

5.11 Monsters

NEWT: My mommy always said there were no monsters. No real ones. But there are.
(James Cameron's *Aliens* (1986))

If one does not want to make the assumptions about the semantics of knowledge ascriptions that were required by the operator account and the lambda abstraction account, there is a further option that deserves serious consideration. That is to treat 'says' as a context shifting operator, an operator that is capable of shifting the context parameter. According to the operator account, 'says' is an operator that shifts the epistemic standard parameter of the index. However, if 'says' were a context shifting operator, we could provide a similar account of the data involving indirect speech reports. Instead of shifting a parameter of the index, we can think of 'says' as shifting a parameter of the context.²⁷ This kind of view would make more or less the same predictions as the operator account.²⁸ When 'knows' is embedded under 'says', the relevant standards will be the standards of the subject.

The existence of such operators is controversial. Kaplan (1989) famously called such operators 'monsters', and claimed that not only are there no such operators in English, but that it is impossible to add such operators to it.

Are there such operators as 'In some contexts it is true that', which when prefixed to a sentence yields a truth if and only if in some context the contained *sentence* (not the content expressed by it) expresses a content that is true in the circumstances of that context?

²⁷See e.g. Anand and Nevins (2004) for a discussion of how to develop the idea that attitude verbs are context shifting operators in a more precise and rigorous way.

²⁸I am not going to discuss whether a view that treats 'says' as a context shifting operator would make exactly the same predictions as a view that treats it as an operator that shifts a parameter of the index. This could for instance depend on the existence of other operators and how they interact with the relevant attitude verbs. However, in the relatively simple cases that I am interested in, this does not matter. Thanks to Derek Ball and Dilip Ninan for relevant and helpful discussion.

Let us try it:

(107) In some contexts it is true that I am tired now.

For [(107)] to be true in the present context it suffices that some agent of some context not be tired at the time of that context. [(107)], so interpreted, has nothing to do with me or the present moment. But this violates Principle 2! Principle 2 can also be expressed in more theory laden way by saying that indexicals always take primary scope. If this is true—and it is—then no operator can control the character of the indexicals within its scope, because they will simply leap out of its scope to the front of the operator. I am not saying that we could not construct a language with such operators, just that English is not one. And such operators *could not be added to it*. (Kaplan, 1989, p. 510, original emphasis)

However, it is not entirely clear what is going on in this passage, and on what grounds is Kaplan denying that there could be monsters in English. If he is offering an empirical argument, which is how Schlenker (2003) appears to interpret him, a generalisation based on a single case is not exactly overwhelming evidence. Furthermore, it is also unclear how much weight we should put on the claim that the existence of monsters conflicts with the principle that indexicals are directly referential.

A thorough discussion of these issues is beyond the scope of the present discussion. However, recent work in linguistics by, among others, Schlenker (2003), Anand and Nevins (2004), and Anand (2006), suggest that there could be monsters in at least some natural languages such as Amharic, Slave, and Zazaki. In particular, it is argued that attitude verbs are capable of shifting indexicals in these languages. Having said that, it still remains controversial whether there are any monsters in English. Schlenker (2003) argues that certain temporal indexicals like ‘two days ago’ can be shifted by attitude verbs, but Anand and Nevins (2004) suggest that the relevant examples involve anaphora, not a genuine context shift.²⁹ In any case, while I do not want to take a stand on the more general questions about the nature and existence of monsters, I think it is safe to say that

²⁹More recently, Santorio (forthcoming) has argued that all indexicals in English are shiftable

the jury is still out when it comes to the questions of whether there is independent evidence that there are monsters in English.

Treating ‘says’ as a monster allows contextualists to explain how ‘knows’ behaves when it is embedded in indirect speech reports without having to posit an epistemic standard parameter of the index or the existence of a covert variable associated with ‘knows’. In particular, the view is compatible with a more traditional understanding of ‘knows’ as an indexical that gets its value directly from the context. However, given the controversial status of monsters, it is a serious commitment if one wants to treat ‘says’ in this way for the purpose of defending a more simple contextualist semantic for ‘knows’. The price one has to pay in order to avoid having to make the assumptions about the semantics of ‘knows’ that were made in the previous sections, is to make some rather controversial assumptions about the semantics of ‘says’. The operator account also required us to make assumptions about the semantics of ‘says’. But in that case we were mainly expanding on the role of ‘says’ as an intensional operator, allowing it to shift the epistemic standard parameter as well as the world parameter. In either case, ‘says’ is shifting a parameter of the index, not a parameter of the context.

5.12 Another Problem

At this point, one might worry that there is a further problem involving indirect speech reports that has not been addressed. Let us consider how it would make sense for John to reason with regards to what Mary said.³⁰ Let us start out by assuming that (85) is true in John’s context.

by attitude verbs and epistemic modals. If that is true, it would offer some support to a contextualist who wants to treat ‘says’ as a monster in order to account for the behaviour of ‘knows’ when it is embedded in indirect speech reports. But even if Santorio is correct to claim that all English indexicals are shiftable, a contextualist would still have to hold that ‘knows’ is shiftable in a different way than ‘I’. After all, an important part of the original problem was that ‘I’ and ‘knows’ behave differently when embedded in indirect speech reports. This does not rule out of an account that treats ‘says’ as a monster, but it means that one has to be careful about what counts as independent evidence for that claim. The mere claim that indexicals are shiftable in some way or another, is not sufficient.

³⁰The way I am presenting this problem is similar to the way that Egan et al. (2005) present a similar problem involving epistemic modals. See e.g. also Hawthorne (2004).

(85) Mary says, 'Smith knows that he has hands'.

So far we have focused on the question of how contextualists can make sense of the inference from (85) to (77).

(77) Mary says that Smith knows that he has hands.

I have argued that there are ways for contextualists to respect the judgement that it is acceptable for John to make this inference. However, there is a further problem. Let us suppose that (108) is also true in John's context.

(108) Smith doesn't know that he has hands.

In that case, one might think that it makes sense for John to conclude that (109) is also true.

(109) What Mary says is false.

The kind of strategies that I have explored on behalf of contextualism, do not deliver this verdict. For instance, according to the operator account, what Mary says is the proposition that Smith knows that he has hands relative to her epistemic standards. Since that proposition is true regardless of what the relevant epistemic standards are, (109) is predicted to be false. Similar considerations apply if we appeal to lambda abstraction or monstrous operators. In either case, the prediction is that (109) is false.

I am not going to address this challenge here, though I will note that I find it less clear that (109) is true as uttered by John, than it is that (77) is a true report of what Mary says.³¹ In any case, my concern is with the original challenge involving indirect speech reports. However, I think it is worth noting that these considerations do not provide any straightforward motivation for nonindexical contextualism either. If we accept a nonindexical contextualist semantics for knowledge ascriptions, (109) would be true in John's context. What Mary says is the proposition that Smith knows that he has hands and that proposition is false relative to the epistemic standards that are operative in John's context. This

³¹Even if it is true in the case of knowledge ascriptions that John can say things like (109), this may not hold for other expressions. For instance, in the case of predicates of taste, it strikes me as inappropriate to say things like 'What she said is false' on this kind of basis.

might seem like a good thing, but there are reasons why one might not want to get too excited about this. Insofar as it is plausible that (109) is true in John's context, one might think that goes for (110) as well.

(110) Mary's utterance of 'Smith knows that he has hands' is false.

But a nonindexical contextualist semantics does not predict that (110) is true. Let us assume the definition of utterance truth given by MacFarlane (2009).

(111) An utterance of a sentence s in a context of use c_U is true relative to a context of assessment c_A if and only if the proposition expressed by s in c_U is true relative to the world of c_U and the epistemic standards of c_A

Given this definition of utterance truth, a nonindexical contextualist semantics for knowledge ascriptions delivers the prediction that (110) is false in John's context. What this means is that nonindexical contextualism can deliver the prediction that (109) is true, but that is at the cost of allowing the truth of what a speaker says, propositional truth, to come apart from the truth of the utterance in question. This is not a welcome consequence. For instance, as Cappelen and Hawthorne (2009, p. 22) have pointed out, it sounds strange to say things like (112).³²

(112) Mary's utterance is true, but what she said is false.

³²MacFarlane (2009) takes utterance truth to be a technical notion, and he argues that we should not put too much weight on judgements involving this notion. He argues that when it comes to ordinary speech, truth is attributed to propositions rather than utterances. If one takes this line, one might be less worried about utterance truth and propositional truth coming apart. However, I am not entirely convinced by the reasons he gives. MacFarlane (2009) follows Kaplan (1989) in taking semantics to be in the business of assigning truth-values to sentences rather than utterances. I am sympathetic to this point of view. However, it is not necessary to believe that utterance truth plays a significant role in semantic theory in order to believe that we have an independent grip on the notion of utterance truth.

Furthermore, even if we grant that truth is typically attributed to propositions rather than utterances in ordinary speech, and that propositional truth is somehow more fundamental than utterance truth, it does not follow that utterance truth is a purely technical notion in the sense that we have no independent grip on it. It still does not seem altogether implausible that we have some grip on this notion and that this is reflected in our judgements about sentences like (112).

If we want to predict that (110) is true as uttered by John, we need to go beyond nonindexical contextualism. According to the relativist treatment of knowledge ascriptions defended by MacFarlane (2005a, forthcoming-b), sentences are true or false relative to not only a context of use, but also a context of assessment. If we adopt the definition of utterance truth in (113), this kind of view allows us to predict that Mary's utterance is false relative to John's context of assessment.

- (113) An utterance of a sentence s in a context of use c_U is true relative to a context of assessment c_A if and only if the proposition expressed by s in c_U is true relative to the world, time, and location of c_U and the epistemic standards of c_A .

Where does this leave us? If the choice comes down to giving up (109) and allowing utterance truth and propositional truth to come apart, it is far from clear that we should prefer the former. That is not to say that this does not constitute a serious challenge for contextualism, only that it is not clear that nonindexical contextualism fares much better. I have not said anything that rules out a relativist view of the sort defended by MacFarlane (2005a, forthcoming-b). There is also the possibility of taking these considerations to motivate some kind of invariantist view about knowledge ascriptions.

While I am not going to try to provide an answer to this challenge on behalf of contextualism here, I think that there are good reasons to separate this challenge from the original challenge. It is the original problem involving indirect speech reports that I have been concerned with in this chapter. The original problem was a problem concerning how speakers can report what another speaker in a different context says or believes. The challenge that has been discussed in this section introduces another set of considerations. It has to do with attributions of truth and falsity, and how speakers evaluate, not just report, what another speaker says in a different context. While these problems are connected, we do not have to go through indirect speech reports in order to generate problems involving attributions of truth and falsity. For instance, we can generate similar problems using propositional anaphora. To the extent that it is appropriate for John to utter (109) when he overhears what Mary says, that is true of (114) as well.

(114) That is false!

Someone who denies that John can use (77) to report what Mary says does not avoid these problems. It still makes sense to attempt to respond to the original problem involving indirect speech reports, even if a solution to that problem does not necessarily amount to a solution to the problem that was presented in this section. I have argued that contextualists can respond to the original problem involving indirect speech reports if they pay sufficient attention to the compositional semantics of knowledge ascriptions and indirect speech reports. What the additional challenge that has been discussed in this section illustrates is that we should not expect these kinds of strategies to solve all the problems for contextualism about knowledge ascriptions.

Conclusion

What is the upshot of all of this? The main conclusion is that there are ways for contextualists to deal with the problem of indirect speech reports. More specifically, it appears that we do not need to adopt a nonindexical contextualist treatment of knowledge ascriptions in order to make sense of the behaviour of knowledge ascriptions that are embedded in indirect speech reports. For instance, according to the operator account, the index contains an epistemic standard parameter and 'says' shifts the value of this parameter to the standards of the subject. This seems to be sufficient to give us the right results in the problematic cases while preserving the spirit of the original contextualist proposal. Along similar lines, the lambda abstraction account treats 'knows' as associated with a variable that can be bound by a lambda operator.

However, in order to get these results we had to make certain assumptions about the semantics of knowledge ascriptions. This suggests that there is a sense in which we can learn something about the semantics of knowledge ascriptions by looking at the way they behave in embedded contexts like indirect speech reports. The data involving indirect speech reports may not give us sufficient reasons for rejecting contextualism altogether, but it may still give us good reasons for implementing contextualism in a certain way. Moreover, the assumptions made by the operator account and the lambda abstraction account can

themselves be subject to further testing. If we want to assess whether a contextualist semantics for knowledge ascriptions is plausible, these are considerations we ought to take into account.

Conclusion

In Chapter 1 I introduced some of the issues having to do with disagreement and relative truth. I considered a simple relativist account of disagreement and discussed an objection to that account. The objection was based on what happens to disagreement if propositional truth is relative to parameters like worlds and times. I considered two lines of response on behalf of relativism. The first response was based on rejecting the premise that propositions are true or false relative to parameters like worlds and times. The second response was based on arguing that the parameters that relativists want to posit are different from parameters like worlds and times. In light of these considerations, the argument against was found to be inconclusive. However, it was not clear to what extent the relativist account of disagreement could be said to be independently motivated.

The focus in Chapter 2 was on more general questions about agreement and disagreement. In particular, I focused on the question of what kinds of attitudes may be involved when two parties agree or disagree. I argued against the so-called ‘doxastic view’, the view that agreement and disagreement is always a matter of what the parties believe. Instead, I developed a more pluralist view of agreement and disagreement by building on Stevenson’s (1937, 1944, 1963) notion of ‘disagreement in attitude’. According to this view, agreement and disagreement can be a matter of what the parties believe, but it may also be a matter of what they want, what they like, and so forth.

The idea that we should not assume a too narrow conception of agreement and disagreement, has played an important role throughout a lot of the discussion. In Chapter 3 I discussed contextualism and relativism about predicates of taste. While contextualism is supposed to have problems with disagreement, I argued that the argument against contextualism is based on a too narrow con-

ception of disagreement. If a sufficiently wide range of cases of disagreement are taken into account, the argument against contextualism looks less convincing. In Chapter 4 I argued for a similar conclusion regarding contextualism about epistemic modals.

While I have argued that the disagreement data do not give us reasons for preferring relativism over contextualism in the case of predicates of taste and epistemic modals, it should be clear that I am not offering anything like a silver bullet that can make all the problems involving agreement and disagreement disappear. If we can draw any general lessons based on this discussion, it is that agreement and disagreement are complicated phenomena. One should be careful about relying on considerations involving agreement and disagreement, but that is not a reason for dismissing these considerations out of hand. We have encountered recalcitrant data, for instance with respect to expressions like ‘probably’, that are difficult to account for. I have also not made any claims about whether the defence of contextualism about predicates of taste and epistemic modals can be straightforwardly extended to other contextualist views. For all I have said, it is still an open question whether contextualism about, say, knowledge ascriptions is compatible with a satisfactory account of disagreement. That question can only be addressed by paying close attention to the relevant disagreement data.

Chapter 5 differed from the other chapters insofar as I did not talk about disagreement. Instead, I focused on another problem for contextualist views, namely the problem of indirect speech reports. Focusing on contextualism about knowledge ascriptions, I argued that there are ways for contextualists to make sense of the behaviour of knowledge ascriptions that are embedded in indirect speech reports. However, in order to do that one must make certain assumptions about the semantics of knowledge ascriptions and indirect speech reports. I also suggested that insofar as these assumptions can be subject to empirical tests, this can serve as the basis for further research. From the point of view of looking for an overall assessment of contextualism, the issues involving indirect speech reports are still relevant.

If what I have said is correct, the resources required to deal with disagreement and the resources required to deal with indirect speech reports are not the same. In order to deal with the problems involving disagreement, it is important to

appreciate what it takes to for two parties to disagree. We need to pay close attention to the attitudes of the parties more generally and not just with regard to what they believe. But when one is dealing with the problems involving indirect speech reports, it becomes important to think about the compositional semantics for the relevant expressions.

Torfinn Thømesen Huvenes

February 2011

References

- Anand, P. (2006). *De De Se*. Unpublished doctoral dissertation, MIT.
- Anand, P., & Nevins, A. (2004). Shifty Operators in Changing Contexts. *Proceedings of SALT*, 14, 20-37.
- Asher, N., & Lascarides, A. (2009). Agreement, Disputes and Commitments in Dialogue. *Journal of Semantics*, 26, 109-158.
- Beaney, M. (Ed.). (1997). *The Frege Reader*. Oxford: Blackwell.
- Blackburn, S. (1984). *Spreading the Word*. Oxford: Clarendon Press.
- Blackburn, S. (1998). *Ruling Passions*. Oxford: Clarendon Press.
- Brogaard, B. (2008). In Defence of a Perspectival Semantics for 'Know'. *Australasian Journal of Philosophy*, 86, 439-459.
- Cappelen, H. (2008a). Content Relativism and Semantic Blindness. In M. García-Carpintero & M. Kölbel (Eds.), *Relative Truth* (p. 265-286). Oxford: Oxford University Press.
- Cappelen, H. (2008b). The Creative Interpreter. *Philosophical Perspectives*, 22, 23-46.
- Cappelen, H., & Hawthorne, J. (2009). *Relativism and Monadic Truth*. Oxford: Oxford University Press.
- Cappelen, H., & Hawthorne, J. (2011). Reply to Glanzberg, Soames and Weatherston. *Analysis*, 71, 143-156.
- Cappelen, H., & Hawthorne, J. (forthcoming). Reply to Lasersohn, MacFarlane, and Richard. *Philosophical Studies*. (to appear)
- Cappelen, H., & Lepore, E. (1997). Varieties of Quotation. *Mind*, 106, 429-450.
- Cappelen, H., & Lepore, E. (2005). *Insensitive Semantics*. Oxford: Blackwell.
- Cohen, S. (1988). How to be a Fallibilist. *Philosophical Perspectives*, 2, 91-123.
- Cohen, S. (1999). Contextualism, Skepticism, and the Structure of Reasons.

- Philosophical Perspectives*, 13, 57-89.
- Cumming, S. (2005). Two Accounts of Indexicals in Mixed Quotation. *Belgian Journal of Linguistics*, 17, 77-88.
- DeRose, K. (1991). Epistemic Possibilities. *Philosophical Review*, 100, 581-605.
- DeRose, K. (1992). Contextualism and Knowledge Attributions. *Philosophy and Phenomenological Research*, 52, 913-929.
- DeRose, K. (2006). Bamboozled by Our Own Words. *Philosophy and Phenomenological Research*, 73, 316-338.
- DeRose, K. (2009). *The Case for Contextualism*. Oxford: Oxford University Press.
- Dietz, R. (2008). Epistemic Modals and Correct Disagreement. In M. García-Carpintero & M. Kölbel (Eds.), *Relative Truth* (p. 239-262). Oxford: Oxford University Press.
- Dreier, J. (1999). Transforming Expressivism. *Noûs*, 33, 558-572.
- Dreier, J. (2009). Relativism (and Expressivism) and the Problem of Disagreement. *Philosophical Perspectives*, 23, 79-110.
- Egan, A. (2007). Epistemic Modals, Relativism and Assertion. *Philosophical Studies*, 133, 1-22.
- Egan, A. (2010). Disputing about Taste. In R. Feldman & T. Warfield (Eds.), *Disagreement* (p. 247-286). Oxford: Oxford University Press.
- Egan, A., Hawthorne, J., & Weatherson, B. (2005). Epistemic Modals in Context. In G. Preyer & G. Peter (Eds.), *Contextualism in Philosophy* (p. 131-168). Oxford: Clarendon Press.
- von Fintel, K., & Gillies, A. (2007). An Opinionated Guide to Epistemic Modality. In T. S. Gendler & J. Hawthorne (Eds.), *Oxford Studies in Epistemology: Volume 2* (p. 32-62). Oxford: Oxford University Press.
- von Fintel, K., & Gillies, A. (2008). CIA Leaks. *Philosophical Review*, 117, 77-98.
- von Fintel, K., & Gillies, A. (forthcoming). *Might Made Right*. In A. Egan & B. Weatherson (Eds.), *Epistemic Modality*. Oxford: Oxford University Press. (to appear)
- von Fintel, K., & Heim, I. (ms.). *Intensional Semantics*. (Unpublished manuscript, MIT)
- Francén, R. (2010). No Deep Disagreement for New Relativists. *Philosophical Studies*, 151, 19-37.
- Frege, G. (1918). Der Gedanke: Eine Logische Untersuchung. *Beiträge zur Philoso-*

- phie des Deutschen Idealismus*, 1, 58–77. (Reprinted in Beaney (1997), p. 323–345.)
- García-Carpintero, M., & Kölbel, M. (Eds.). (2008). *Relative Truth*. Oxford: Oxford University Press.
- Geurts, B. (1998). The Mechanisms of Denial. *Language*, 74, 274–307.
- Gibbard, A. (1990). *Wise Choices, Apt Feelings*. Oxford: Clarendon Press.
- Gibbard, A. (2003). *Thinking How to Live*. Cambridge, MA: Harvard University Press.
- Glanzberg, M. (2007). Context, Content, and Relativism. *Philosophical Studies*, 136, 1–29.
- Grice, H. P. (1989a). Logic and Conversation. In *Studies in the Way of Words* (p. 22–40). Cambridge, MA: Harvard University Press.
- Grice, H. P. (1989b). *Studies in the Way of Words*. Cambridge, MA: Harvard University Press.
- Hacking, I. (1967). Possibility. *Philosophical Review*, 76, 143–168.
- Hawthorne, J. (2004). *Knowledge and Lotteries*. Oxford: Clarendon Press.
- Hawthorne, J. (2006). Testing for Context-Dependence. *Philosophy and Phenomenological Research*, 73, 443–450.
- Hawthorne, J. (2007). Eavesdroppers and Epistemic Modals. *Philosophical Issues*, 17, 92–101.
- Heim, I., & Kratzer, A. (1998). *Semantics in Generative Grammar*. Oxford: Blackwell.
- Humberstone, L. (2006). Sufficiency and Excess. *Proceedings of the Aristotelian Society*, 80, 265–320.
- Huvenes, T. (forthcoming). Varieties of Disagreement and Predicates of Taste. *Australasian Journal of Philosophy*. (to appear)
- Ichikawa, J. (forthcoming). Quantifiers and Epistemic Contextualism. *Philosophical Studies*. (to appear)
- Jackson, F. (2008). The Argument from the Persistence of Moral Disagreement. In R. Shafer-Landau (Ed.), *Oxford Studies in Metaethics: Volume 3* (p. 75–86). Oxford: Oxford University Press.
- Jackson, F., & Pettit, P. (1998). A Problem for Expressivism. *Analysis*, 58, 239–251.
- Joyce, J. (2010). A Defense of Imprecise Credences in Inference and Decision Making. *Philosophical Perspectives*, 24, 281–323.

- Kaplan, D. (1989). Demonstratives. In J. Almog, J. Perry, & H. Wettstein (Eds.), *Themes from Kaplan* (p. 481-563). Oxford: Oxford University Press.
- King, J. C. (2003). Tense, Modality, and Semantic Values. *Philosophical Perspectives*, 17, 195-245.
- Kölbel, M. (2002). *Truth Without Objectivity*. London: Routledge.
- Kölbel, M. (2004). Faultless Disagreement. *Proceedings of the Aristotelian Society*, 104, 53-73.
- Kolodny, N., & MacFarlane, J. (2010). Ifs and Oughts. *Journal of Philosophy*, 107, 115-143.
- Kratzer, A. (1977). What *Must* and *Can* Must and Can Mean. *Linguistics and Philosophy*, 1, 337-355.
- Kratzer, A. (1981). The Notional Category of Modality. In H.-J. Eikmeyer & H. Rieser (Eds.), *Words, Worlds, and Contexts* (p. 38-74). Berlin: de Gruyter.
- Kratzer, A. (1991a). Conditionals. In A. von Stechow & D. Wunderlich (Eds.), *Semantik/Semantics: International Handbook of Contemporary Research* (p. 651-656). Berlin: de Gruyter.
- Kratzer, A. (1991b). Modality. In A. von Stechow & D. Wunderlich (Eds.), *Semantik/Semantics: International Handbook of Contemporary Research* (p. 639-650). Berlin: de Gruyter.
- Laserson, P. (2005). Context Dependence, Disagreement, and Predicates of Personal Taste. *Linguistics and Philosophy*, 28, 643-686.
- Lewis, D. (1978). Truth in Fiction. *American Philosophical Quarterly*, 15, 37-46. (Reprinted in Lewis (1983), p. 261-275.)
- Lewis, D. (1979). Attitudes De Dicto and De Se. *Philosophical Review*, 88, 513-43. (Reprinted in Lewis (1983), p. 133-155.)
- Lewis, D. (1980). Index, Context, and Content. In S. Kanger & S. Öhman (Eds.), *Philosophy and Grammar* (p. 79-100). Dordrecht: Reidel. (Reprinted in Lewis (1998), p. 21-44.)
- Lewis, D. (1983). *Philosophical Papers* (Vol. 1). Oxford: Oxford University Press.
- Lewis, D. (1986). *On the Plurality of Worlds*. Oxford: Blackwell.
- Lewis, D. (1996). Elusive Knowledge. *Australasian Journal of Philosophy*, 74, 549-567. (Reprinted in Lewis (1999), p. 418-445)
- Lewis, D. (1998). *Papers in Philosophical Logic*. Cambridge: Cambridge University Press.

- Lewis, D. (1999). *Papers in Metaphysics and Epistemology*. Cambridge: Cambridge University Press.
- Ludlow, P. (2005). Contextualism and the New Linguistic Turn in Epistemology. In G. Preyer & G. Peter (Eds.), *Contextualism in Philosophy* (p. 11-50). Oxford: Clarendon Press.
- Ludlow, P. (2008). Cheap Contextualism. *Philosophical Issues*, 18, 104-129.
- MacFarlane, J. (2003). Future Contingents and Relative Truth. *Philosophical Quarterly*, 53, 321-336.
- MacFarlane, J. (2005a). The Assessment Sensitivity of Knowledge Attributions. In T. S. Gendler & J. Hawthorne (Eds.), *Oxford Studies in Epistemology: Volume 1* (p. 197-233). Oxford: Oxford University Press.
- MacFarlane, J. (2005b). Making Sense of Relative Truth. *Proceedings of the Aristotelian Society*, 104, 53-73.
- MacFarlane, J. (2007). Relativism and Disagreement. *Philosophical Studies*, 132, 17-31.
- MacFarlane, J. (2008). Truth in the Garden of Forking Paths. In M. García-Carpintero & M. Kölbel (Eds.), *Relative Truth* (p. 81-102). Oxford: Oxford University Press.
- MacFarlane, J. (2009). Nonindexical Contextualism. *Synthese*, 166, 231-250.
- MacFarlane, J. (2011). What is Assertion? In J. Brown & H. Cappelen (Eds.), *Assertion: New Philosophical Essays* (p. 79-96). Oxford: Oxford University Press.
- MacFarlane, J. (forthcoming-a). Epistemic Modals are Assessment-Sensitive. In A. Egan & B. Weatherson (Eds.), *Epistemic Modality*. Oxford: Oxford University Press. (to appear)
- MacFarlane, J. (forthcoming-b). Relativism and Knowledge Attributions. In S. Bernecker & D. Pritchard (Eds.), *Routledge Companion to Epistemology*. London: Routledge. (to appear)
- MacFarlane, J. (ms.). Varieties of Disagreement. (Unpublished manuscript, University of California, Berkeley)
- Maier, E., & van Der Sandt, R. (2003). Denial in Discourse. (Unpublished manuscript, University of Nijmegen)
- Manley, D. (2009). A Guided Tour of Metametaphysics. In D. Chalmers, D. Manley, & R. Wasserman (Eds.), *Metametaphysics: New Essays on the Foundations*

- of Ontology*. Oxford: Oxford University Press.
- Maudlin, T. (2007). Review of Patrick Greenough and Michael Lynch (Eds.), *Truth and Realism*. *Notre Dame Philosophical Reviews*. (URL: <http://ndpr.nd.edu/review.cfm?id=10143>)
- Moruzzi, S. (2008). Assertion, Belief and Disagreement: A Problem for Truth-Relativism. In M. García-Carpintero & M. Kölbel (Eds.), *Relative Truth* (p. 207-224). Oxford: Oxford University Press.
- Nagel, J. (2010). Knowledge Ascriptions and the Psychological Consequences of Thinking About Error. *Philosophical Quarterly*, 60, 286-306.
- Ninan, D. (2010). De Se Attitudes: Ascription and Communication. *Philosophy Compass*, 5, 551-567.
- Ninan, D. (forthcoming). Semantics and the Objects of Assertion. *Linguistics and Philosophy*. (to appear)
- Nunberg, G. (1993). Indexicality and Deixis. *Linguistics and Philosophy*, 16, 1-43.
- Papafragou, A. (2006). Epistemic Modality and Truth Conditions. *Lingua*, 116, 1688-1702.
- Portner, P. (2009). *Modality*. Oxford: Oxford University Press.
- Predelli, S. (2005). *Contexts*. Oxford: Clarendon Press.
- Preyer, G., & Peter, G. (Eds.). (2005). *Contextualism in Philosophy*. Oxford: Clarendon Press.
- Recanati, F. (2007). *Perspectival Thought*. Oxford: Oxford University Press.
- Richard, M. (1981). Temporalism and Eternalism. *Philosophical Studies*, 39, 1-13.
- Richard, M. (2004). Contextualism and Relativism. *Philosophical Studies*, 119, 215-242.
- Richard, M. (2008). *When Truth Gives Out*. Oxford: Oxford University Press.
- Rosenkranz, S. (2008). Frege, Relativism and Faultless Disagreement. In M. García-Carpintero & M. Kölbel (Eds.), *Relative Truth* (p. 225-237). Oxford: Oxford University Press.
- Santorio, P. (forthcoming). Modals are Monsters. *Proceedings of SALT*. (to appear)
- Schaffer, J. (2004). From Contextualism to Contrastivism. *Philosophical Studies*, 119, 73-104.
- Schaffer, J. (forthcoming-a). Contextualism for Taste Claims and Epistemic Modals. In A. Egan & B. Weatherson (Eds.), *Epistemic Modality*. Oxford: Oxford University Press. (to appear)

- Schaffer, J. (forthcoming-b). *Necessitarian Propositions*. *Synthese*. (to appear)
- Schlenker, P. (2003). A Plea for Monsters. *Linguistics and Philosophy*, 26, 29-120.
- Schroeder, M. (2008). *Being For*. Oxford: Clarendon Press.
- Simons, M. (2007). Observations on Embedding Verbs, Evidentiality, and Pre-supposition. *Lingua*, 117, 1034-1056.
- Stalnaker, R. (1978). Assertion. In P. Cole (Ed.), *Syntax and Semantics 9* (p. 315-332). New York: Academic Press. (Reprinted in Stalnaker (1999), p. 78-95.)
- Stalnaker, R. (1999). *Context and Content*. Oxford: Oxford University Press.
- Stanley, J. (2000). Context and Logical Form. *Linguistics and Philosophy*, 23, 391-434.
- Stanley, J. (2005). *Knowledge and Practical Interests*. Oxford: Clarendon Press.
- Stanley, J., & Szabó, Z. (2000). On Quantifier Domain Restriction. *Mind and Language*, 15, 219-261.
- Stephenson, T. (2007). Judge Dependence, Epistemic Modals, and Predicates of Personal Taste. *Linguistics and Philosophy*, 30, 487-525.
- Stevenson, C. L. (1937). The Emotive Meaning of Ethical Terms. *Mind*, 46, 14-31.
- Stevenson, C. L. (1944). *Ethics and Language*. New Haven: Yale University Press.
- Stevenson, C. L. (1963). The Nature of Ethical Disagreement. In *Facts and Values* (p. 1-9). New Haven: Yale University Press.
- Stojanovic, I. (2007). Talking about Taste: Disagreement, Implicit Arguments, and Relative Truth. *Linguistics and Philosophy*, 30, 691-706.
- Swanson, E. (2006). *Interactions with Context*. Unpublished doctoral dissertation, MIT.
- Weatherson, B. (2008). Attitudes and Relativism. *Philosophical Perspectives*, 22, 527-544.
- Weatherson, B. (2009). Conditionals and Indexical Relativism. *Synthese*, 166, 333-357.
- Weatherson, B. (2011). No Royal Road to Relativism. *Analysis*, 71, 133-143.
- White, R. (2010). Evidential Symmetry and Mushy Credences. In T. S. Gendler & J. Hawthorne (Eds.), *Oxford Studies in Epistemology: Volume 3* (p. 161-186). Oxford: Oxford University Press.
- Williamson, T. (2000). *Knowledge and its Limits*. Oxford: Oxford University Press.

- Wright, C. (2007). New Age Relativism and Epistemic Possibility: The Question of Evidence. *Philosophical Issues*, 17, 262-283.
- Yalcin, S. (2007). Epistemic Modals. *Mind*, 116, 983-1026.
- Yalcin, S. (2010). Probability Operators. *Philosophy Compass*, 5, 916-937.
- Yalcin, S. (forthcoming). Nonfactualism about Epistemic Modality. In A. Egan & B. Weatherson (Eds.), *Epistemic Modality*. Oxford: Oxford University Press. (to appear)